Does the Consolidated Feed Matter?

Shihao Yu*

Current version: June 16, 2022

Abstract

In this paper we examine the role of the consolidated feeds in the current U.S. equities market by studying exogenous events which affect their speed and availability respectively. We find that faster consolidated feeds have an adverse, albeit mild, impact on market liquidity, possibly as a result of more non-high-frequency algorithmic trading activities from informed institutional traders. Moreover, when the consolidated feeds become corrupted or unavailable due to technical glitches, market liquidity significantly worsens. Our findings suggest that the consolidated feeds remain a crucial component of today's market data infrastructure.

^{*}School of Business and Economics, Vrije Universiteit Amsterdam. Email: s.yu@vu.nl. I thank Albert J. Menkveld, Bart Yueshen Zhou, Björn Hagströmer for their helpful suggestions and discussions. In addition, I thank participants at the Microstructure Exchange online seminar series for their helpful comments. I gratefully acknowledge Netherlands Organization for Scientific Research (NWO) for a Research Talent grant.

1 Introduction

Trading in the U.S. equities market has become increasingly complicated in the past two decades. First of all, trading is highly fragmented across 16 public exchanges, dozens of broker-dealer internalizers and dark pools. Moreover, trading is extremely fast at sub-millisecond frequencies.¹ In such a fragmented and fast trading environment, having access to low-latency market data is crucial for various market participants to implement their trading strategies: market makers need fast data to constantly re-price their quotes to avoid adverse selection; institutional investors need it to find the best available prices to execute their trades and arbitragers need it to exploit short-lived arbitrage opportunities.

There are two main types of market data in today's U.S. equities market: consolidated feeds disseminated by what's called security information processors (SIPs) and direct feeds disseminated by exchanges. While the consolidated feeds or the SIP feeds² are mandated by regulation, direct feeds are expensive, proprietary products of the exchanges and normally have a lower latency and contain more information than consolidated feeds.³ So some market participants argue that the discrepancy between the two types of feeds has created an unfair "two-tiered" market with the haves (direct feeds subscribers) and have not (consolidated feeds subscribers). Perhaps based on such reasoning, the SEC adopted a new rule named Reg NMS II in February 2021, aiming to further improve the consolidated feeds and make it more comparable to direct feeds and thus level the playing field of market data. However, the rule has caused a fierce backlash from exchanges⁴ which claim that further enhancements to the consolidated feeds will not benefit the market.⁵ Across the Atlantic, ESMA is scheduled to roll out a European-wide consolidated tape in the foreseeable future. Details over what should be included, e.g., only post-trades or including pre-trade quotes, and what the optimal reporting latency is still very much in debate.

Objective. One argument by the exchanges that enhancing the consolidated feeds will not

¹For example, Menkveld (2018) analyzes a sample of Nasdaq trades in October 2010 and finds that twenty percent of trades arrive in sub-millisecond clusters.

²In what below, we will use the consolidated feeds and the SIP feeds interchangeably.

³See Section 2 below for details about the speed and content difference between the consolidated feeds and direct feeds.

⁴SEC approved the rule in December, 2020 but was sued by Nasdaq, NYSE and Cboe. Thus the implementation of the rule is now blocked. See "Nasdaq, NYSE Sue SEC to Block Market Data Overhaul.", *The Wall Street Journal*, February 9, 2021.

⁵"Building the SIP Autobahn", Nasdaq, May 20, 2021.

benefit the market is that it is primarily used by "display users" or human traders. Thus, further speed improvement in the sub-second frequencies seems unlikely to benefit them. As for latency-sensitive traders, they will always opt for faster direct feeds and be hardly affected either. However, anecdotal evidence shows that the consolidated feeds are more used by algorithmic traders as it becomes faster.⁶ Moreover, there are scenarios where the consolidated feeds are useful to market participants. For example, regulatory filings show that about 45% dark pools do not use any direct feeds but the consolidated feeds instead.⁷ In addition, even for low-latency traders, they normally use data feeds from multiple sources for data integrity check (CFTC and SEC, 2010) and might withdraw from their trading when the consolidated feeds become unreliable (Aldrich, Grundfest, and Laughlin, 2017). A recent paper by Ernst, Sokobin, and Spatt (2021) shows that fast traders react to off-exchange trade reports from the consolidated feeds even though they appear with a significant delay. So the key empirical question is: do the consolidated feeds really matter in today's market, and if so, in what way? Our objective is to try to shed some light on the question.

Approach. To identify the role of the consolidated feeds, we exploit events which exogenously affect their speed and availability respectively. The first event we use is a technology upgrade to the SIP operated by Nasdaq on October 24, 2016, which significantly reduced its *processing* latency: the median drops from about 350 microseconds to less than 20 microseconds. In the U.S. equities market, there are two SIPs operated by Nasdaq and NYSE respectively, with the former being responsible for disseminating consolidated feeds for Nasdaq-listed stocks (Tape C securities) and the latter for NYSE-listed stocks (Tape A securities) and stocks listed on other regional exchanges and their successors (Tape B securities). Such a unique structure of two SIPs allows us to perform a standard difference-in-difference (DiD) analysis based on a matched sample of Nasdaq-listed and NYSE-listed stocks. If the speed of the consolidated feeds matters for market participants, we should see changes in liquidity and trading in Nasdaq-listed stocks after the upgrade, *relative to* NYSE-listed stocks.

In addition, we exploit the unique geography of exchanges and two SIPs for further identification. Specifically, while Nasdaq exchanges and Nasdaq-SIP are both located at Carteret, New Jersey, NYSE exchanges are about 35 miles away at Mahwah, New Jersey (See Figure 1). Thus, for

⁶"Consolidated Market Data Feeds Gain Traction in Algo Trading and Fixed Income", Finextra, January 2019.

⁷See, for example, "Dispelling the Complementary Product Theory for Market Data", Nasdaq, August 8, 2020.

Nasdaq-listed stocks traded on Nasdaq exchanges, traveling time of a message (e.g., trade or quote update) to the SIP is small and therefore processing time at the SIP makes up the largest component of total SIP latency. As the Nasdaq-SIP upgrade slashes its processing latency by over 90%, for Nasdaq-listed stocks, total SIP latency of a message from Nasdaq exchanges has a much larger relative reduction than that from NYSE exchanges (See Section 3.3.1 for a detailed explanation). So we expect the differential impact on Nasdaq-listed stocks relative to NYSE-listed stocks, if it exists, to be more pronounced on Nasdaq exchanges versus other exchanges.

Then we turn to events where SIPs experienced technical glitches and the consolidated feeds became corrupt or unavailable. Given only one of the two SIPs is affected during all our sample events, we are able to use similar DiD identification strategy as above based on a matched sample of stocks with their consolidated feeds affected and ones unaffected. To examine market liquidity during SIP glitch events, we obtain direct feed data from MayStreet to compute high-frequency measures as consolidated feeds such as NYSE TAQ database is by definition not available. In addition, we exploit a unique feature of the Nasdaq-SIP glitch event on January 3, 2013 for further identification. Specifically, the glitch first occurred in even-numbered data dissemination channels ("early channels") of the Nasdaq-SIP and only occurred several minutes later in odd-numbered channels ("late channels"). So during the first glitch period, stocks allocated to the late channels can serve as an ideal control group and allows for a clean DiD analysis.

Findings. Our major findings can be summarized as follows. First, we find that faster SIP feeds lead to a mild worsening of market liquidity. A difference-in-difference analysis shows that, after the Nasdaq-SIP upgrade which significantly reduces the SIP latency for Nasdaq-listed stocks, their quoted spread and price impact increase *relative to* NYSE-listed stocks. In addition, we show that the upgrade gives rise to an increase in overall algorithmic trading (AT) activity measured by quote-to-trade ratio, but not high-frequency trading (HFT) activity measured by strategic runs (Hasbrouck and Saar, 2013). So we conjecture that the worsening of market liquidity might result from an elevated level of non-HFT AT activity from informed institutional investors. Second, we document that when SIP feeds are corrupted or unavailable due to technical glitches, market liquidity deteriorates significantly, especially in terms of market trading volume and order-book depth. Taking stock, our findings show that the consolidated feeds play an important role in

today's market data infrastructure and past speed improvements might have led to an unexpected and negative impact on market liquidity. Thus, future proposals for change should be carefully assessed.⁸

Literature. Our paper relates to several strands of literature. First, it is directly related to literature studying the effect of market data on trading and market quality. Brogaard, Ringgenberg, and Rösch (2020) examine events where exchanges start to charge a fee for their proprietary feeds for the first time. They find that the introduction of data fees leads to a significant fall in the market volume of the fee charging exchange and it is mainly due to it having less time at the NBBO and getting less inter-market sweep (ISO) orders. Hendershott, Rysman, and Schwabe (2020) study the event when NYSE introduces its new data product, NYSE Integrated Feed, and show that there is a complementary relationship between exchange's proprietary data sales and trading activity: firms increased their share of trading on NYSE after the introduction. My paper differs from them in that we focus on the consolidated feed instead of exchanges' proprietary feeds. By examining events that exogenously affect the speed and availability of the consolidated feed, we show that the consolidated feed matters in today's market and shapes trader behavior and exchange competition, contributing to the ongoing debate over the market data reform. Ye, Yao, and Gai (2013) examine an older speed upgrade to the Nasdaq-SIP and find it has no impact on overall market liquidity. However, they do not study its impact on trader behavior and competition.

Second, our paper adds to the literature on the difference between consolidated feeds and direct feeds. O'Hara, Yao, and Ye (2014) show that odd-lot trades were missing in the consolidated feeds⁹ which account for a large share of trading volume and are quite informed. They conjecture that odd-lot trader are used by the informed traders to hide their trading intention from the consolidated feeds Battalio, Corwin, and Jennings (2016) analyze a sample of high-priced stocks and show that the exclusion of odd-lot orders from consolidated feeds results in some trades being filled at worse prices. Ding, Hanna, and Hendershott (2014) compare NBBO constructed from the consolidated feeds and NBBO constructed by adding direct feeds from Nasdaq, BATS and Direct Edge exchanges. They find that the dislocation between the two NBBOs can happen quite frequently for active stocks

⁸It should be noted that our analysis cannot directly speak to the potential impact of the SEC's newly proposed data overhaul plan by, e.g., adding more information such as depth to the consolidated feed and having competing consolidators.

⁹Odd-lot trades were added into the consolidated feed after a SIP reform in 2013.

but its duration is quite short. So the cost for low-frequent traders is small. Similarly, Bartlett and McCrary (2019) use the exchange timestamp included in the consolidated feeds to construct NBBO with zero-latency and show that profit from direct feed arbitrage is not economically significant. Hasbrouck (2019) compares the price discovery contribution of SIP compared to theoretically constructed direct feeds and finds the latter dominates at high frequency. My paper contributes to the existing studies by examining the event when the speed advantage of direct feeds over consolidated feeds largely narrows. Compared with the approach used in the existing studies that statically compares the two feeds, my paper looks at a real-life event so that it can incorporate the effect of traders' dynamic responses. However, we acknowledge that we only focus on events of two specific types, i.e., SIP's speed upgrade and glitch. So we cannot speak to other currently proposed changes to the consolidated feed such as including depth information, odd-lot quotes and auction imbalances.

Third, our paper closely relates to studies that examine the impact of trading speed on market quality¹⁰. Based on the adverse selection channel, theoretical models have shown that the impact of trading speed is ultimately dependent on which trader groups become faster. Market liquidity worsens if short-term informed arbitrageurs become faster (Biais, Foucault, and Moinas, 2015; Budish, Cramton, and Shim, 2015; Foucault, Hombert, and Roşu, 2016), improves if market makers become faster (Hoffmann, 2014; Jovanovic and Menkveld, 2016), and depends on market conditions (e.g., news arrival frequency and the presence of noise traders) when both trader groups become faster (Menkveld and Zoican, 2016). Empirical studies largely support the theoretical predictions above. For example, Brogaard, Hagströmer, Nordén, and Riordan (2015) show that market liquidity improves when market-makers opt for a speed upgrade of their co-location server to the exchange. Shkilko and Sokolov (2020) find that market liquidity improves when the microwave network between Chicago and New York is disrupted by weather conditions, which makes high-frequency arbitrageurs slower. My paper adds to the literature by examining the event of a speed upgrade to the consolidated feed, which increases the trading speed (as least the market data component) of slow traders (buy-side execution algorithms) who do not have access to direct feeds due to cost considerations or technical complexities. We find the speed upgrade has no impact on overall market liquidity, suggesting that a faster consolidated feed might have equally affected arbitrageurs

¹⁰See Menkveld (2016) for a comprehensive review on this topic.

and market makers. However, the arm race between them seems to intensify as low-latency trading activity increases after the upgrade. It might be due to both fast market makers and arbitrageurs responding to a faster trading speed of slow traders.

2 Institutional background

Before detailing our data and identification strategy, we first provide an overview of the institutional background regarding the market data structure of the U.S. equities market.

2.1 A two-tiered market data structure

Broadly speaking, there are two types of market data in today's U.S. equities market. First, there are the consolidated feeds mandated by the SEC. The consolidated feeds are disseminated by the Security Information Processors (SIPs), which, on a near real-time basis, collect trades and top-of-book quote updates from all national securities exchanges such as NYSE and Nasdaq, aggregate them into consolidated tapes and disseminate them to their subscribers. In addition to trades and quote updates, the SIPs disseminate critical regulatory information including the National Best Bid and Offer (NBBO), Limit Up-Limit Down (LULD) price bands, short sale restrictions, and halts. Subscribers of the SIPs form a heterogeneous group, including TV and media, broker-dealers, investment advisors, and algorithmic traders, and they are charged by varying rates depending on their specific type. Revenue of the SIPs are then shared across all participating exchanges based on their contribution to market trading volume and depth at the NBBO.¹¹

Second, there are direct feeds sold by exchanges as their proprietary data products. Compared with the consolidated feeds, direct feeds are faster and contain more information such as depthof-book information, odd-lot quotes and auction imbalance information. As a result, sophisticated traders such as high-frequency traders use direct feeds as key inputs to their trading algorithms. However, direct feeds are prohibitively expensive for unsophisticated traders. A 2019 report published by IEX on market data cost estimates that the total annual subscription fee to all national securities exchanges' direct feeds sum up to around 1.15 million US dollars, excluding costs for

¹¹See "SIP Accounting 101", Nasdaq, March 25, 2021, at https://www.nasdaq.com/articles/sip-accounting-101-2021-03-25

physical and logical connectivity needed to receive the feeds (IEX, 2019).

Such a two-tiered market data structure has spurred heated debates over whether it is unfair to small market participants as they in a way suffer from information asymmetry due to inferior market data. The SEC has been improving the consolidated feeds to bring them closer to direct feeds in terms of speed and content. Besides the processing latency reduction, which is the focus of the paper, the most recent proposal by SEC called NMS 2.0 plans to add depth-of-book information (up to five levels), odd-lot quotes, auction imbalance information to consolidated feeds.¹²

2.2 Unique features of the consolidated feeds

Two SIPs The first salient feature of the consolidated feeds for U.S. equities is that they are not disseminated by one SIP, but two separate ones. Specifically, one SIP is managed by NYSE ("NYSE-SIP") and responsible for disseminating consolidated feeds for securities listed on NYSE (Tape A Securities) or other regional exchanges and their successors (Tape B Securities, e.g., ETPs listed on NYSE ARCA). Instead, the other SIP is managed by NASDAQ ("NASDAQ-SIP") and responsible for disseminating consolidated feeds for securities listed on NASDAQ (Tape C Securities). So the bottomline is that there is a separation between where a security is traded and where the trade is reported: trades and quotes of a stock from any national exchange have to be reported to the NYSE-SIP as long as its listing exchange is NYSE, and to the NASDAQ-SIP if its listing exchange is NASDAQ.

Geography of SIPs Another important feature of SIPs relates to their unique geography. As depicted in Figure 1, all major US equities exchanges and the two SIPs are located in a triangular area in New Jersey, with (1) Nasdaq venues (Nasdaq, BX, and PSX) and the Nasdaq-SIP at Carteret, NJ, (2) NYSE venues (NYSE, NYSE Arca, NYSE American) and the NYSE-SIP in Mahwah and (3) Cboe venues (BZX, BYX, EDGX and EDGA) in Secaucus. The geography of SIPs and exchanges creates what's called "geographical latency" of the consolidated feeds. For example, a trade or quote update for a Nasdaq stock from NYSE venues will have to first travel from Mahwah to Careret to get processed by the Nasdaq-SIP. Note that NMS 2.0, the new market data proposal by

Figure 1. The New Jersey "Equity Triangle". This figure sketches the geographical locations of major US equity exchanges and two SIPs. Besides, it shows the estimated traveling latencies between three data centers and processing latencies at the two SIPs before the upgrade event on October 24, 2016.¹³



SEC, plans to create multiple SIPs, aiming exactly to solve the issue of the geographical latency.

2.3 SIP upgrade events

The two SIPs have received several technology upgrades over the recent years, resulting in significant reductions in their processing latency. As Figure 2 shows, there were three major upgrades: (1) NASDAQ-SIP on October 24, 2016, (2) NYSE-SIP at the end of 2018 and (3) NYSE-SIP on July 13, 2020. In our analysis below we focus on the Nasdaq-SIP upgrade on October 24, 2016 as the latency reduction is most significant: its median processing latency drops more than 90% from about 350 microseconds to less than 20 microseconds. Nasdaq implements the technology upgrade of its SIP by migrating it to the Nasdaq Financial Framework and INET, the proprietary technology behind Nasdaq exchanges. As a result, not only was there a significant decrease in the Nasdaq-SIP's processing latency, but also an increase in its resiliency & reliability, capacity, scalability and message efficiencies.¹⁴

¹²See details of the proposed rule at https://www.sec.gov/news/press-release/2020-311

¹⁴See "Nasdaq Sets the Record Straight About the SIP", Nasdaq, October 25, 2016, https://www.nasdaq.com/articles/nasdaq-sets-record-straight-about-sip-2016-10-25.

Figure 2. SIP speed upgrades. Starting from August 2015, each TAQ message has two timestamps: Participant Timestamp (when the message is registered at the exchange from which it originates) and SIP Timestamp (when the message is disseminated by the SIP). Thus we can compute the SIP latency by the difference between the two timestamps. To compute the processing latency of the NYSE-SIP, we use quotes in General Electric originating from the NYSE exchange so that there is minimal traveling latency. By the same token, to compute the processing latency of the Nasdaq-SIP, we use quotes in Apple originating from the Nasdaq exchange. The figures plots the daily median latency for NYSE-SIP and Nasdaq-SIP respectively.



2.4 SIP glitch events

[] While SIP glitches are not rare, there are only a few with market-wide impact. Going through

 Table 1. Recent market-wide SIP glitches. This table lists some key information about the three market-wide SIP glitch events that happened in recent years.

Date	Start and End Time	Duration	SIP	Market-wide Trading halt
January 3, 2013	13:33 - 13:51	18 minutes	Nasdaq-SIP	No ^a
October 30, 2014	13:07 - 13:34	27 minutes ^b	NYSE-SIP	No ^c
August 12, 2019	15:15 - 15:27	12 minutes ^b	NYSE-SIP	No

^a There is no market-wide trading halt. EDGX and EDGA halted trading for Nasdaq-listed stocks after 13:42.

^b In both events, The NYSE shifted operations to its disaster recovery site in Chicago after the glitch was solved.

^c Some dark pools, including ITG Posit and Goldman Sachs' Sigma X, which uses the NYSE SIP were closed during the glitch period.

all system alerts published by the two SIPs¹⁵, we find three market-wide SIP glitch events in recent years: (1) the Nasdaq-SIP glitch on January 3, 2013, (2) the NYSE-SIP glitch on October 30, 2014, and (3) the NYSE-SIP glitch on August 12, 2019. In Table 1, we provide some basic information about the three SIP glitches such as the start and end time, duration and whether there was a

¹⁵Market data alerts for the NYSE-SIP and Nasdaq-SIP can be found at https://www.ctaplan.com/alerts# and https://www.utpplan.com/vendor_alerts respectively.

trading halt during the glitch period. Then we offer a succinct account for each SIP glitch event

afterwards.

Table 2. Channel assignment of the NASDAQ-SIP and outage order. This table shows the symbol allocation across the six data dessimination channels of the Nasdaq-SIP. Moreover, it shows the starting and ending time of the glitch for trades and quotes in each channel.

Outage order	Outage order Channel		Trade outage period	
"Late"channels	Channel 1 (Symbols A-CDZ) Channel 3 (Symbols FE-LKZ) Channel 5 (Symbols PC-SPZ)	13:37:22 - 13:48:19	13:36:51 - 13:51:14	
"Early" channels	Channel 2 (Symbols CE-FDZ) Channel 4 (Symbols LL-PBZ) Channel 6 (Symbols SQ-ZZZ)	13:33:11 - 13:48:21	13:33:11 - 13:51:15	

Nasdaq-SIP glitch on January 3, 2013 As summarized in Table 2, on January 3, 2013, at 13:33:11 ET, the dissemination network of the Nasdaq-SIP lost connectivity, causing its even-numbered channels to cease dissemination of both trades and quote updates. A few minutes later, the remaining odd-numbered channels also ceased dissemination of trades and quote updates at 13:36:51 and 13:37:22 respectively. To present visual evidence of the glitch, Figure A3 plots the trade and quote counts from the SIP feeds and direct feeds for stocks by dissemination channel type. It shows that the number of SIP trades and quotes quickly dropped to zero for stocks in the early channels when the first glitch hit and then stocks in the late channels when the second glitch hit. In contrast, direct feeds operated normally during the whole glitch period.

NYSE-SIP glitch on October 30, 2014 On October 30, 2014, at approximately 13:07 ET, the NYSE-SIP was hit by a hardware failure which impacts its data feed dissemination. Figure A4 in the appendix plots trade and quote counts from the SIP feeds and direct feeds for a sample of NYSE-listed stocks and Nasdaq-listed stocks respectively. It shows that at the start of the glitch, trade and quote count of NYSE-listed stocks quickly dropped to almost zero. Note that for treated stocks, not SIP trades and quotes from all exchanges are missing. A closer examination of the SIP feeds shows that, for NYSE-listed stocks, it is either trades and quotes from Nasdaq or Bats that are missing.

At about 13:34 ET, the NYSE-SIP was switched to its backup data center in Chicago.¹⁶

NYSE-SIP glitch on August 12, 2019 On August 12, 2019, at approximately 15:15 ET, the NYSE-SIP experienced a hardware failure of one of its network core routers, causing disruptions to the dissemination of both trades and quote updates. Again, Figure A5 in the appendix plots trade and quote counts of the SIP feeds versus direct feeds for a matched sample of NYSE-listed stocks and Nasdaq-listed stocks. It shows that, during the glitch period, for NYSE-listed stocks only quote updates from NYSE Arca appear in the SIP feeds and trades from all exchanges are missing. As the event on October 30, 2014, the operation of NYSE-SIP was later switch to the backup data center in Chicago at approximately 15:27 ET.¹⁷.

We would like to mention that the three SIP glitch events share two important features. First, unlike the Nasdaq "Flash Freeze" event on August 22, 2013 which resulted in a three-hour trading halt for Nasdaq-listed stocks, there was no market-wide trading halt during any of the three SIP glitch event. Second, while the consolidated feeds were unreliable or unavailable during glitches, direct feeds from exchanges operated normally. Although the duration of the three SIP glitch events is fairly short and their economic impact might be small, they provide us an opportunity to empirically study what happens to the market when an importance piece of data infrastructure is missing and there is maximum information asymmetry: sophisticated traders such as HFTs who do not reply on the consolidated feeds but subscribe to direct feeds from exchanges are not directly affected; in contrast, retail brokers and dark pools which rely on the consolidated feeds will receive unreliable or no pricing information.

3 Data and Identification Strategy

3.1 Data sources

¹⁶https://www.bloomberg.com/news/articles/2014-10-30/disaster-averted-in-nyse-stocks-as-backup-feed-kicksin

¹⁷https://www.ctaplan.com/alerts#110000144324

3.1.1 Nasdaq-SIP upgrade on October 24, 2016

To assess the overall impact of the Nasdaq-SIP upgrade on market functioning, we examine a wide-range of liquidity and trading variables. To do so, we combine several datasets from different sources. The first dataset is the NYSE TAQ database¹⁸, which is essentially NYSE's version of the SIP feeds and includes trades and top-of-book quote updates from all national securities exchanges. With the NYSE TAQ, we are able to calculate common liquidity and trading variables such as bid-ask spread, top-of-book depth and trading volume. The second dataset we use is the SEC's MIDAS¹⁹, which collects direct feeds from all national securities exchange combination, MIDAS reports daily count and volume of new order submissions, count of cancel messages, hidden volume, and odd-lot volume. The third and last dataset is the LOBSTER data²⁰, an adapted copy of the Nasdaq ITCH dataset, which contains full order-book event messages such as new order submissions, cancellations and trade executions.

Our sample covers a matched sample of 296 Nasdaq-listed stocks and the same number of NYSE-listed stocks over the period from August 29, 2016 to December 16, 2016 (See Section 3.3.1 for details about our matching procedures.). So the sample period spans a four-month window, two months before and two months after the event date of October 24, 2016. A four-month window length is chosen to strike a balance between a too-long window which might include other events and a too-short window which might not generate sufficient statistical power. As a robustness check, in Section A.1 of the appendix, we shorten the window to two month, one month before and one month after the event date of October 24, 2016, and the results do not change significantly.

3.1.2 SIP glitch events

When analyzing SIP glitch events, we use instead direct feeds data collected by MayStreet, a US data company and supplier of SEC's MIDAS. Direct feeds from different exchanges have different formats.²¹ For some direct feeds such as NYSE OpenBookUltra, they contain only level book mes-

¹⁸The NYSE TAQ dataset is accessed through WRDS (Wharton Research Data Services).

¹⁹https://www.sec.gov/marketstructure/downloads.html

²⁰https://lobsterdata.com/index.php

²¹Specifically, we use the following direct feeds: BATS BZX Multicast Pitch, BATS BYZ Multicast Pitch, CHX Book Feed, DirectEdge EDGA Multicast EdgeBook Depth, DirectEdge EDGX EDGX Multicast EdgeBook Depth, NASDAQ

sages which represent an update on a single level of the order book. In other words, we can not see individual order submissions and cancellation as all orders at a particular price are communicated in a consolidated format. In contrast, direct feeds from other exchanges are similar to the LOBSTER data and contain full order-book event messages, i.e., new order submissions, cancellations and executions. Either case, we can use the direct feed messages from each exchange to build the exchange-specific limit order book. Then we obtain the "consolidated" limit order book for the whole market by aggregating limit order books of all exchanges. We use the consolidated limit order book to compute several market-wide and high-frequency liquidity and trading measures. It is perhaps worth noting that the direct feeds from exchanges are necessary in order to study the SIP glitch events as, by definition, SIP feeds (such as the NYSE TAQ) is either unreliable or unavailable during such events. As with the analysis of the Nasdaq-SIP upgrade, we include a matched sample of Nasdaq-listed stocks and NYSE-listed stocks. In addition, for each SIP glitch event, we include observations in the glitch period, the "event-window" and a half-an-hour window before the start of the glitch, the "pre-event window" (See Table 1 for detailed timelines for each glitch event).

3.2 Liquidity and trading variables

We first summarize in Table 3 all liquidity and trading variables used in the analysis of the SIP upgrade and SIP glitch events respectively. Then we provide a detail account of each.

3.2.1 Nasdaq-SIP upgrade

NBBO based on Participant Timestamps Since August 2015, SIP trade and quote messages start to include two timestamps: a what's called Participant Timestamp, the time when the message is registered at the originating exchange, and a SIP Timestamp, the time when the message is disseminated by the SIP. As argued above in Section 2 on institutional details, SIP trade and quote messages can be subject to delays due to either processing latency at the SIP or traveling latency from the originating exchange to the SIP. Thus the SIP Timestamp will be later than corresponding Participant Timestamp and the NBBO constructed based on the SIP Timestamp ("SIP-NBBO") will

BX TotalView-ITCH, NASDAQ TotalView-ITCH, NASDAQ PSX TotalView-ITCH, NYSE MKT OpenBook Ultra, NYSE ARCA ARCABook, ARCA Trades, NYSE OpenBook Ultra, NYSE Trades, National Stock Exchange Multicast Depth of Book.

Table 3. Liquidity and trading variables. This table summarizes the liquidity and trading variables used in the analysis of the SIP upgrade and SIP glitch events respectively. It reports the names, scope, definition and data sources for each variable.

Event	Measure	Scope	Definition	Data Source
	RQS	Market-wide & Exchange-specific	Relative quoted spread	NYSE TAQ
	RES	Market-wide & Exchange-specific	Relative effective spread	NYSE TAQ
	RRS	Market-wide & Exchange-specific	Relative realized spread	NYSE TAQ
	Depth	Market-wide & Exchange-specific	Depth at NBBO	NYSE TAQ
	Vlm	Market-wide & Exchange-specific	Trading volume	NYSE TAQ
SIP-upgrade	PrcImp	Market-wide	Retail trade price improve- ment	NYSE TAQ
	ISOShr	Market-wide & Exchange-specific	Share of ISO trades	NYSE TAQ
	OddlotShr	Market-wide & Exchange-specific	Share of odd-lot trades	NYSE TAQ
	Cancel/Trade	Market-wide & Exchange-specific	Ratio of # cancels to # trades	SEC MIDAS
	Order/Trade	Market-wide & Exchange-specific	Ratio of add limit order vol- ume to trading volume	SEC MIDAS
	#Run/Vlm	Nasdaq	Ratio of # strategic run to trading volume	LOBSTER
	RQS	Market-wide	Relative quoted spread	MayStreet
	RES	Market-wide	Relative effective spread	MayStreet
SID alitab	RRS	Market-wide	Relative realized spread	MayStreet
SII-gilleli	Vlm	Market-wide	Trading volume	MayStreet
	DepthNBBO	Market-wide	Depth at NBBO	MayStreet
	Depth5Lvl	Market-wide	Depth cumulative across five best bids and asks	MayStreet

lag that based on the Participant Timestamp ("Participant-NBBO").

The discrepancy between the SIP-NBBO and Participant-NBBO can lead to biases in certain liquidity measures if they are computed based on the SIP-NBBO included in the NYSE TAQ database. For example, trade-related liquidity measures such as effective spread require one to first infer the direction of a trade, i.e., whether it is a buyer- or seller initiated. One popular trade classification algorithm is Lee and Ready (1991): trades with a transaction price higher (lower) than the prevailing midquote will be classified as buy (sell) trades. So latency embedded in the SIP-NBBO can result in misalignment between trades and their actual prevailing NBBOs and thus wrong classifications. Another measure that might be biased based on SIP-NBBO is one exchange's NBBO depth contribution. Imagine there are two exchanges, Exchange A and Exchange B, and the former always sets the new NBBO first and the latter follows. However, due to geographical latency, Exchange B's quote updates might be processed by the SIP first before Exchange A's arrive. Thus Exchange B will always be the exchange which sets the new NBBO. To make things more

complicated, SIP processing time is significantly reduced in the SIP-upgrade event we study, which might mechanically change the liquidity measures. Given the foregoing reasons, we construct the BBOs for all exchanges and then the NBBO *based on the Participant Timestamp*. In addition, when matching trades with their prevailing NBBOs, we use trades' Participant Timestamp as well.

Liquidity variables We use several measures capturing two different aspects of market liquidity. The first three captures the price aspect of liquidity: relative quoted spread (*RQS*), relative effective spread (*RES*) and relative realized spread (*RRS*). Relative quoted spread, RQS, measures the cost of a round trip of small trades and is defined as:

$$RQS_t = \frac{NBO_t - NBB_t}{Mid_t},$$
(1)

where *t* is the timestamp of the current snapshot of the order book. NBO_t and NBB_t are the national best offer and bid respectively. Mid_t is the midquote, which is simply $(NBO_t + NBB_t)/2$.

However, RQS can be a poor proxy for the actual transaction cost traders pay for two reasons. First, for a large trade which "walks up the order book", i.e., is executed against quotes at multiple price levels, its average transaction price is worse than the prevailing NBBO. Second, it is common for trades to be executed at prices better than NBBO: off-exchange trades can either be executed at the NBBO midquote in dark pools or receive price improvement from wholesalers like Citadel and Virtu; on-exchange trades, instead, can be executed against hidden orders priced better than NBBO. Given all the concerns above, relative effective spread, *RES* can better measure the actual transaction cost and is defined below:²²

$$\operatorname{RES}_{t} = \frac{d_{t} \left(p_{t} - \operatorname{Mid}_{t} \right)}{\operatorname{Mid}_{t}},$$
(2)

where *t* indexes trades. d_t is the trade direction indicator. p_t is the actual transaction price for trade *t*.²³ Mid_t is the prevailing midquote *just before* the trade.

²²When computing trade-based liquidity measures such as effective spread, I drop all regular trades which include Stock-Option Trade, Average Price Trade, Derivatively Priced Trade and etc. These trades often have execution prices far off from the prevailing market price, skewing their trade-based liquidity measures. Thus I drop those trades when computing the daily metrics. Specifically, I only keep trade records with trade conditions "@", "F", "I", "F I", for Tape A securities and "@F", "I", "@F I", "@ I", for Tape C securities. Moreover, when deciding the trade direction, we use the Exchange Timestamp instead of the SIP Timestamp.

²³For trades with multiple executions, we use the volume weighted average price.

Another related liquidity measure is relative realized spread, RRS, defined below:

$$RRS_t = \frac{d_t \left(p_t - Mid_{t+\Delta_t} \right)}{Mid_t},$$
(3)

where $\operatorname{Mid}_{t+\Delta_t}$ is the prevailing midquote after a time interval of Δ_t after the trade. I pick the common choice of 30 seconds. The relative realized spread is essentially the relative effective spread less the price impact and is a crude proxy for the profits of market makers²⁴.

In addition to the liquidity measures above, we compute two other measures capturing the quantity aspect of liquidity : dollar depth at NBBO (*Depth*), dollar trading volume (*Vlm*). A deeper dollar depth at NBBO makes large trades cheaper. While trading volume depends on various factors including volatility and information. But given the same volatility and information level, better market liquidity leads to larger trading volume as traders find it cheaper to trade and realize their private values.

Besides, we in particular examine the execution quality of retail trades, which are most likely benchmarked by retail brokers against the SIP-NBBO and thus might fall victim to its embedded latency.²⁵ Following Boehmer, Jones, Zhang, and Zhang (2021), we label a trade as a retail trade if it is executed at a sub-penny price, resulting from a price improvement received from a wholesaler like Citadel and Virtu. For example, for a trade executed at 10.011 (10.009), it will be labeled as a retail sell (buy) trade with a price improvement of 0.1 cent per share. Then we are able to compute, at the stock-day level, share-size weighted average price improvements of retail trades and use it as a proxy for their execution quality (*PrcImp*). Importantly, we benchmark execution prices of retail trades with their prevailing Participant-NBBOs as opposed to SIP-NBBOs so that we can see whether they receive "true" price improvement. For example, a broker/wholesaler can claim that a retail trade receives a price improvement, but it is based on a stale SIP-NBBO.

²⁴The relative realized spread is only a crude proxy for market maker profit. For example, it does not include the rebates market makers receive from the exchanges and various costs from co-location, exchange data subscription and fixed IT costs.

²⁵Anecdotal evidence shows that wholesalers might use stale SIP quotes to price their client trades. On January 13, 2017, the US SEC fined Citadel Securities \$22.6 million dollars for the use of two algorithms that "did not internalize retail orders at the best price observed nor sought to obtain the best price in the marketplace." Citadel's high frequency trading strategy exploited difference between prices on SIP and the more accurate direct exchange feeds.

Trading variables In addition to liquidity, we compute several trading variables to capture the trading behavior of market participants. The first two are the trading volume via inter-market sweep order (ISO) as a fraction of total trading volume (*ISOShr*) and odd-lot trading volume as a fraction of total trading volume (*OddlotShr*).²⁶ Then we compute two proxies for algorithmic trading: the ratio of cancel order count to trade count (*Cancel/Trade*) and the ratio of volume of add limit orders to trading volume (*Order/Trade*). Hagströmer and Nordén (2013) show that market-making HFTs have substantially higher quote-to-trade ratios than other fast traders as they constantly reprice their quotes in order to avoid adverse selection. So an increase in the quote-to-trade ratio is likely to reflect an increase in market-making related activities. However, the drawback of such quote-to-trade measures is that they also capture activities of non-HFTs such as execution algorithms (EAs), which might seek passive execution and thus frequently reprice their quotes in order book. For example, Beason and Wahal (2021) analyze a large dataset of 2.3 million parent orders executed by institutional investors and find less than 0.4% of their child orders are market orders.

To better measure low-latency HFT activities, we follow Hasbrouck and Saar (2013) and identify long strategic runs from the Nasdaq order book event messages provided by LOBSTER. A strategic run is a series of linked submissions, cancellations, and executions, representing dynamic order placement strategies.²⁷ As Hasbrouck and Saar (2013) argues, while the strategic run measure might reflect activities from agency algorithms, it is more likely that *long* strategic runs predominately captures HFT activities. In fact, empirical studies find the measure to be highly correlated with true HFT activity measures, both time-series and cross-section (Hasbrouck and Saar, 2013; Yao and Ye, 2018). Moreover, following Yao and Ye (2018), we scale the raw number of strategic runs by trading volume (#*Run/Vlm*).²⁸

All liquidity measures are first computed at the tick-by-tick frequency and later aggregated

²⁶ISO trades and odd-lot trades are identified in the TAQ with a trade indicator of "F" and "I" respectively.

²⁷Specifically, to identify one strategic run, we start with a new limit order submission and link it with its subsequent cancellation or execution based on the same reference number provided by the exchange. If it is a cancellation, then we check whether it is followed by either a subsequent new limit order submission or an execution within 100 ms in the same direction and for the same size. In addition, if a limit order is partially executed, and the remainder is cancelled, we check whether a subsequent re-submission or execution based on the cancelled quantity.

²⁸Note that except for the price improvement and strategic run measure, all other liquidity and trading variables above are computed both for the entire market aggregated across all exchanges and for each exchange separately. For example, *Depth* of an exchange will be its own dollar depth at the NBBO.

Table 4. Summary statistics: the Nasdaq-SIP upgrade event on October 24. 2016. *RQS, RES, RPI,* and *RRS* stand for relative quoted spread, relative effective spread, relative price impact and relative realized spread respectively and are all in basis point. *Depth* is NBBO depth in thousand dollars. *PrcImp* is the share-size weighted average price improvements received by the retail trades and in cents per hundred shares. *Vlm* is trading volume in million dollars. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *ISOShare* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. *#Run/Vlm* is the number of strategic runs per million dollar trading volume. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. The sample period is from August 29 to December 16, 2016.

Variable	Ν	Mean	SD	Min	50%	Max
RQS (bp)	42900	8.55	9.01	0.85	6.19	310.27
RES (bp)	42900	2.65	3.15	0.28	1.85	204.06
RPI (bp)	42900	2.35	2.08	-12.02	1.83	100.72
RRS (bp)	42900	0.29	2.12	-44.76	0.01	128.58
Depth (\$ thousand)	42900	160.11	505.77	6.94	76.59	14241.81
Vlm (\$ million)	42900	104.23	238.14	0.20	46.24	11120.27
PrcImp (¢/100Shr)	42900	14.55	3.31	1.00	14.68	34.71
ISOShr (%)	42900	35.03	7.23	3.66	34.96	89.36
OddlotShr (%)	42900	10.75	6.72	0.06	9.65	56.24
Cancel/Trade	42900	23.79	11.98	3.84	21.14	252.44
Order/Trade	42900	38.04	23.76	5.48	32.31	443.62
#Run/Vlm	42900	14.77	14.46	0.00	10.70	519.25

to a given frequency (e.g., daily for the SIP upgrade events). We aggregate stock variables such as *RQS* and *DepthNBBO* by computing their time-weighted averages and flow variables such as *RES* by their dollar-volume weighted averages. Table 4 reports the summary statistics of the above liquidity and trading metrics for the matched sample of Nasdaq-listed and NYSE-listed stocks around the Nasdaq-SIP upgrade event (see Section 3.3.1 below for details on our matching procedure).

3.2.2 SIP glitches

For the analysis of SIP glitch events, we focus on their impact on market-wide liquidity and compute several common liquidity measures. Specifically, we compute the relative quoted spread (*RQS*), relative effective spread (*RES*) and relative realized spread (*RRS*). In addition, we compute the dollar trading volume (*Vlm*), top-of-book depth (*DepthNBBO*) and cumulative depth across five best prices of the order book (*Depth5Lvl*). Note that as we obtain direct feeds from all exchanges during the SIP glitches, we can not only compute the top-of-book depth, but also depth across several levels of the order book. Table 5a and Table 5b report the summary statistics for all liquidity

variables used in the regression for pooling all three sip glitches and for the nasdaq-sip glitch on

January 3, 2013 respectively.

Table 5. Summary statistics: SIP glitch events. This table reports the summary statistics for all liquidity variables used in the regression for pooling all three sip glitches and for the nasdaq-sip glitch on January 3, 2013 respectively. *RQS*, *RES* and *RRS* stand for relative quoted spread, effective spread and realized spread respectively and are all in basis point. *Vlm* is dollar volume in thousand dollars. *DepthNBBO* is NBBO depth in thousand dollars. *Depth5Lvl* is cumulative depth across five best price levels of the order book. The summary statistics are computed base on time-series averages over the sample period.

	Ν	Mean	SD	Min	50%	Max
RQS (bp)	235564	13.80	18.45	0.52	9.08	730.64
RES (bp)	129309	3.32	4.81	0.00	2.04	319.86
RRS (bp)	129309	1.38	6.62	-366.81	0.94	323.56
Vlm (\$ million)	235564	33.30	201.72	0.00	1.82	45054.15
DepthNBBO (\$ thousand)	235564	84.50	347.56	0.07	32.52	34608.91
Depth5Lvl (\$ thousand)	235564	470.26	1387.41	4.25	149.43	42915.56

(a) Three SIP glitch events pooled.

(b) SIP glitch event on January 3, 2013.

Variable	Ν	Mean	SD	Min	50%	Max
RQS (bp)	80886	10.87	9.66	0.88	8.11	140.60
RES (bp)	21803	2.87	3.14	0.00	1.89	54.18
RRS (bp)	21803	1.61	4.15	-48.50	1.30	45.16
Vlm (\$ million)	80886	8.35	49.19	0.00	0.00	2878.37
DepthNBBO (\$ thousand)	80886	126.78	344.84	0.19	37.12	5985.11
Depth5Lvl (\$ thousand)	80886	707.78	1796.72	11.97	196.50	23383.85

3.3 Identification strategy

3.3.1 Nasdaq-SIP upgrade

Thanks to the technology upgrade of the Nasdaq-SIP on October 24, 2016, its processing latency decreased significantly while NYSE-SIP's barely changed (See Figure 2). As the Nasdaq-SIP disseminates the consolidated feeds for Nasdaq-listed stocks while the NYSE-SIP for NYSE-listed stocks, the Nasdaq-SIP upgrade event allows for a clean difference-in-difference (DiD) analysis: after the Nasdaq-SIP upgrade, consolidated feeds become much faster for Nasdaq-listed stocks but not for NYSE-listed stocks. If faster consolidated feeds have any impact on liquidity or trading, we should observe our proxies changes for Nasdaq-listed stocks *relative to* NYSE-listed stocks.

Table 6. Propensity score matching: SIP speed upgrade. This table reports results from the propensity score matching for the SIP speed upgrade. The treatment group consists of 296 NASDAQ-listed stocks are matched with 296 NYSE-listed stocks on price, trading volume, market capitalization and Fama and French 12 industry classification. I use one-to-one nearest neighbor propensity score matching (PSM), without replacement.

		Ν	Mean	SD	10%	25%	50%	75%	90%
Variable	Sample								
Drice (¢)	Control	296	62.03	54.25	14.67	25.52	47.14	79.53	130.23
Price (\$)	Treatment	296	67.35	110.38	10.01	22.08	45.14	79.68	119.85
	Control	296	15.32	26.43	1.93	3.35	6.07	15.24	33.27
MarketCap (& billion)	Treatment	296	16.95	51.23	1.28	3.08	4.93	11.57	30.06
DollarVolume (¢ million)	Control	296	109.34	160.25	15.92	28.75	64.12	124.99	229.15
Donarvolume (\$ million)	Treatment	296	114.08	241.66	13.74	24.99	48.36	102.09	242.71
PSM Score	Control	296	0.41	0.15	0.25	0.29	0.46	0.51	0.61
1 5141 50010	Treatment	296	0.44	0.16	0.25	0.30	0.48	0.55	0.62

To implement the DiD analysis, we construct a matched sample of Nasdaq-listed and NYSElisted stocks. To start with, we follow Brogaard, Ringgenberg, and Rösch (2020) and include all common stocks²⁹ listed on either NYSE or Nasdaq and exclude stocks with dual class shares and a market capitalization below \$500 million. More importantly, we exclude stocks involved in the SEC's Tick Size Pilot Program to avoid its confounding effect. The program started in October 2016 and was conducted by the SEC to assess the impact of wider tick sizes on the liquidity and trading of certain small-capitalization companies ("Pilot Securities"). Pilot Securities are divided into one control group and three test groups. While tick sizes of stocks in the test group remain at \$0.01, those in the test groups increase from \$0.01 to \$0.05 either for their trading or quoting or both.³⁰ After excluding all Pilot Securities, including the control group, we are left with 296 Nasdaq-listed stocks and 633 NYSE-listed stocks.

Then we match the 296 Nasdaq-listed stocks with the same number of NYSE-listed stocks on price, trading volume, market capitalization and industry. The first three matching variables are the daily averages during the month before the event period, that is, between August 1, 2016 and August 23, 2016. In addition, as companies in certain industries have a preference for listing either on Nasdaq or NYSE (e.g., technology companies are more likely to list on Nasdaq), we follow Brogaard, Ringgenberg, and Rösch (2020) and add Fama and French 12 industry classification as a further matching variable.³¹ Moreover, we use one-to-one nearest neighbor propensity score

²⁹Common stocks have a CRSP share code of 10 or 11.

³⁰We refer readers to the program's official website for details (https://www.sec.gov/ticksizepilot)

³¹Perhaps it is worth noting that they use Fama and French 48 industry classification. The reason for me to use the

matching (PSM), without replacement. Table 6 reports the matching results and shows that matching is successful as all three matching variables and the propensity score of the two samples have quite similar support.

The DiD identification approach requires the standard parallel trends assumption, which means the treatment group (i.e., Nasdaq-listed stocks) would have evolved in a similar fashion to the control group (i.e., NYSE-listed stocks) if there was no speed upgrade to the Nasdaq-SIP. Figure A1 in the appendix plots time-series of the liquidity and trading metrics around the Nasdaq-SIP upgrade event and visual evidence suggests that it is indeed the case. In other words, the parallel trends assumption is supported.

Nasdaq-listed vs. NYSE-listed stocks Based on the matched sample described above, I estimate the difference-in-difference (DiD) regressions as follows:

$$metric_{i,t} = \alpha_i + \beta A fter_t + \gamma A fter_t \times Nasdaq Stock_i + \epsilon_{i,t}$$
(4)

where *metrc*_{*i*,*t*} is the liquidity or trading variable of stock *i* on day *t*. α_i is the stock fixed effects. *NasdaqStock*_{*i*,*t*} is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *After*_{*i*,*t*} is dummy variable that equals one on and after October 24, 2016, and equals zero otherwise. Standard errors are clustered at the stock level.

Nasdaq venue vs. other venues The second identification strategy utilizes the unique geography of SIPs and exchanges in the U.S. equities market (See Figure 1). For a trader who subscribes to the SIP feeds, the total latency of receiving a message from one exchange consists of two parts: traveling latency of the message, both from the exchange to SIP and then from SIP to the trader, and processing latency of the message at the SIP. As an example, consider trading in Nasdaq-listed stocks and thus messages from all exchanges have to be reported to and processed at Nasdaq-SIP. For a trader (labeled as "Trader A") who is located at Carteret, NJ where Nasdaq venues are located, the total latency of SIP will be dominated by the processing latency due to small traveling latency from Nasdaq venues to Nasdaq-SIP and from Nasdaq-SIP to the trader as they are all located at

simpler version is due to the relatively small size of the sample. Not every industry has stocks from both tapes in my sample, which makes the logit regression not converge.

the same geographical location. In contrast, for a trader (labeled as "Trader B") who is located at Mahwah, NJ (where NYSE venues are located), a message from NYSE venues (at Mahwah, NJ) has to first travel to Nasdaq-SIP (at Carteret, NJ), be processed there, and then travels back from Nasdaq-SIP (at Carteret, NJ) to the trader (at Mahwah, NJ).

So after the Nasdaq-SIP speed upgrade which reduces the processing latency of the Nasdaq-SIP from 350 microseconds to less than 20 microseconds, the total SIP latency for Trader A is only about 20 microseconds. While for Trader B, the total latency will be 20 microseconds of processing latency plus a round trip traveling latency between Carteret, NJ and Mahwah NJ, which is roughly 560 microseconds through optical fiber and not affected by the reduction of the Nasdaq-SIP processing latency. In other words, the speed up of Nasdaq-SIP will affect Trader A more than Trader B and thus trading on Nasdaq venues more than that on NYSE venues.³²

To implement the identification strategy based on geographical latency, we estimate the following regression specification:

$$metric_{i,t,e} = \alpha_{i,e} + \beta A fter_t + \gamma_1 A fter_t \times NasdaqStock_i + \gamma_2 A fter_t \times NasdaqVenue_e + \gamma_3 A fter_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e}$$
(5)

where *metric*_{*i*,*t*,*e*} is the liquidity or trading variable for stock *i*, traded on exchange *e* and on day *t*. $\alpha_{i,e}$ controls for the stock-venue fixed effects, which is is potentially important as a stock can have a fixed trading pattern on a particular exchange. For example, although Nasdaq-listed stocks can trade at any of the 16 exchanges, opening and closing auctions are only held at Nasdaq, the listing exchange. Besides, for NYSE-listed stocks, NYSE has DMMs (Designated Market Makers) who have mild obligations in maintaining liquidity. Last, exchanges have different fee schedules for stocks listed on its venue compared to stocks otherwise.³³ *After*_{*i*,*t*} is dummy variable that

³²For the sake of the example, here we assume that the two traders, Trader A and Trader B, are located at Carteret, NJ, and Mahwah, NJ, respectively so that the former experiences least geographical latency while the latter the most. However, in reality traders' locations vary. If both traders are located at Secaucus, NJ, for example, the impact of the reduction of the Nasdaq-SIP processing latency will differ less for them due to traveling latency of receiving the SIP feeds.

³³See, e.g., https://www.nasdaqtrader.com/Trader.aspx?id=PriceListTrading2 for the fee schedule of Nasdaq venues.

equals one after October 24, 2016, and equals zero otherwise. *NasdaqStock_i*, is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *NasdaqVenue_i* is a dummy variable that equals one if the variable is computed specifically on exchange *e*, and equals zero otherwise.³⁴ So the coefficient γ_3 captures the triple-difference-in-difference effect, that is, the change in the cross-venue difference for Nasdaq-listed stocks versus NYSE-listed stocks after the Nasdaq-SIP speed upgrade. Standard errors are clustered at the stock level. In Section A.2, as a robustness check, we use an alternative specification where we control for the stock-date fixed effects, instead of stock-venue fixed effects and the results are qualitatively the same.

Again, note that for the DiD identification above to be valid, one needs the parallel assumption to hold on the exchange level, that is, liquidity and trading metrics would have evolved similarly on the Nasdaq exchange relative to other exchanges if there was no speed upgrade to the Nasdaq-SIP. Figure A2 in the appendix plots the time-series of the liquidity and trading metrics series on several exchanges and visual evidence supports the parallel trends assumption.

3.3.2 SIP glitch events

Pooling all three SIP glitch events The identification strategy for the analysis of SIP glitch events is the same as above and uses the unique structure of two SIPs in the U.S. equities market. For each SIP glitch event, it is either the Nasdaq-SIP or the NYSE-SIP that experienced a technical glitch, not both. So when NYSE-SIP experienced a glitch, we expect liquidity and trading in NYSE-listed stocks to be *more* affected than that in Nasdaq-listed stocks and vice versa when Nasdaq-SIP experienced a glitch. Importantly, as Figure A3, A4 and A5 in the appendix show, during all three SIP glitch events, direct feeds were largely unaffected and there is no market-wide trading halt. So the identified effect is not due to trade disruption as in the Nasdaq "Flash Freeze" event on August 22, 2013.

For each SIP glitch event, we use the same matching approach as above to construct a matched sample of 200 treated stocks (e.g., Nasdaq-listed stocks when Nasdaq-SIP experiences a glitch) and the same number of control stocks (e.g., NYSE-listed stocks when Nasdaq-SIP experiences a

³⁴I only include four exchanges NYSE-Arca, Nasdaq, BZX and EDGX as they all adopt a maker-taker model. Moreover, NYSE is excluded as it has only started to trade Nasdaq-listed stocks since April 2018.

Table 7. Propensity score matching: SIP glitch events This table reports results from the propensity score matching for the three SIP glitch events. For each event, the treatment group consists of 1200 randomly chosen stocks whose consolidated feeds are affected by the SIP glitch and are matched with 1200 unaffected stocks on price, trading volume, market capitalization and Fama and French 12 industry classification. I use one-to-one nearest neighbor propensity score matching (PSM), without replacement.

		Ν	Mean	SD	10%	25%	50%	75%	90%
Variable	Sample								
Price (¢)	Control	1200	49.02	63.62	12.15	21.16	34.66	57.20	92.17
The (\$)	Treatment	1200	46.87	81.55	11.59	18.47	31.16	50.06	86.91
MarlatCar (@hillian)	Control	1200	8.42	20.10	0.76	1.36	2.77	7.00	18.07
MarketCap (& Dimon)	Treatment	1200	6.47	34.86	0.60	0.80	1.42	3.56	9.11
DollarVolume (& million)	Control	1200	69.48	132.87	4.14	9.86	25.81	71.78	168.24
Donai volume (\$ minori)	Treatment	1200	60.06	277.91	2.63	4.93	13.24	39.49	113.75
PSM Score	Control	1200	0.53	0.15	0.28	0.44	0.57	0.65	0.70
	Treatment	1200	0.54	0.16	0.28	0.44	0.60	0.68	0.71

glitch). Table 7 reports the propensity matching score results and shows the matching is quite successful with similar distribution of the two sample stocks across the matching variables and the final propensity score. The sample period for each event covers from 30 minutes before the start of the glitch until the end of the glitch. Pooling all three SIP-glitch events, we run the following standard DiD regression below:

$$metric_{i,d,t} = \alpha_{i,d} + \beta A fter_{i,d,t} + \gamma A fter_{d,t} \times Treated_{i,d} + \epsilon_{i,d,t}.$$
(6)

where $metrc_{i,d,t}$ is the liquidity or trading metric of stock *i* on event day *d* during the 30-second time interval *t*. $\alpha_{i,d}$ is the stock-day(event) fixed effect. *Treated*_{*i,d*} is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock on January 3, 2013 and NYSE-listed stock on October 30, 2014 and August 12, 2019. *After*_{*d,t*} is a dummy variable that equals one after the glitch starts, and equals zero otherwise. Standard errors are clustered at the stock-day(event) level. It is perhaps worth noting that the three SIP glitch events happened at different intraday periods and on both the Nasdaq-SIP and NYSE-SIP. Hence, pooling all three events in the regression helps alleviate some concerns from a possibly imperfect matching.

Zoom in onto the event on January 3, 2013 In addition to pooling all three SIP glitch events as above, we zoom in onto the Nasdaq-SIP glitch event on January 3, 2013 and exploit its unique feature for a further identification. Specifically, there are six channels through which the Nasdaq-SIP

disseminates its data feeds and Nasdaq-listed stocks are allocated into the six channels *alphabet-ically*.³⁵ More importantly, as shown in Table 2, the Nasdaq-SIP glitch on January 3, 2013 first occurred at three even-numbered channels and occurred only a few minutes later at the other three odd-numbered channels. Thus stocks in the "late" three channels can serve as a ideal control group during the period between the start of the first and second glitch. Based on the same matching procedure as above, we construct a sampled of 200 randomly chosen Nasdaq-listed stocks belonging to the "early channel", the same number of Nasdaq-listed stocks from the "late channels" and 400 NYSE-listed stocks. In terms of the sample period, we focus on the time interval between 30-minutes before the first glitch and the start of the second glitch, i.e., between 13:03:00 and 13:36:51. Last, we run the following DiD regression:

$$metric_{i,t} = \alpha_i + \beta Period_{i,t} + \gamma_1 Period_{i,t} \times Early Channel_{i,t} + \gamma_2 Period_{i,t} \times Late Channel_{i,t} + \epsilon_{i,t}$$
(7)

where *metrc*_{*i*,*t*} is the liquidity or trading variable of stock *i* in time interval *t*. α_i is the stock fixed effects. *Period*1_{*i*,*t*} is a dummy variable that equals one after the start of the first glitch at 13:33:11, and equals zero otherwise. *EarlyChannel*_{*i*,*t*} is dummy variable that equals one if stock *i* belongs to the early channels, and equals zero otherwise. *LateChannel*_{*i*,*t*} is a dummy variable that equals one if stock *i* belongs to the stock *i* belongs to the late channels, and equals zero otherwise. Note for the regression above, we include stocks from all channels (early, late and normal channels). We expect γ_1 to be significant while γ_2 insignificant if the first glitch only affects stocks in the early channels. Standard errors are clustered at the stock level.

4 **Results**

We next take the identification strategies developed in the previous section to the data and examine the role of consolidated feeds in the current US equities market. We first look at the event of a speed upgrade to the consolidated feeds and then turn to events where the consolidated feeds

³⁵Ye (2012) uses the same feature to study the potential quote stuffing behavior of HFTs and find that messages of stocks within the same channels are more correlated. The allocation of symbols in each of the six channels are according to the alphabetical order. Although the allocation rule might be such that the total message volume in each channel show be more or less the same so that no channel will be suffering constantly high message volume and having higher latency

Figure 3. SIP speed upgrade and NBBO dislocations. Following Bartlett and McCrary (2019), we construct two versions of NBBOs, one based on the SIP Timestamps (SIP-NBBO) and the other based on the Participant Timestamp (Participant-NBBO) and then identify the dislocations between the two NBBOs. By the same token, we construct two BBOs on the Nasdaq exchange, one based again on SIP Timestamp and the other on Participant Timestamp.

(a) This figure plots the median NBBO dislocation duration and number of NBBO dislocations for NYSE-Stocks and Nasdaq-stocks respectively. The time series plotted is the cross-section average of all sample stocks.



(b) This figure plots the median Nasdaq-BBO dislocation duration and number of Nasdaq-BBO dislocations for NYSE-Stocks and Nasdaq-stocks respectively. The time series plotted is the cross-section average of all sample stocks.



experience a technical glitch and thus become unreliable or unavailable.

4.1 SIP speed upgrade

4.1.1 NBBO/Nasdaq-BBO dislocations

Before examining the impact of the SIP upgrade on market liquidity and trading, we first look at to what extend it makes SIP prices more reliable compared with direct feeds. To do so, we follow Bartlett and McCrary (2019) and construct two versions of NBBO, one based on the SIP Timestamp ("SIP-NBBO") and the other based on the Participant Timestamp³⁶ ("Participant-NBBO"). Dislocations between the two NBBOs occur when the SIP-NBBO lags and deviates from the prevailing Participant-NBBO.³⁷ After the Nasdaq-SIP upgrade, which significantly reduces the processing latency of the SIP feed for Nasdaq-listed stocks, we expect them to experience fewer and shorter NBBO dislocations *relative to* NYSE-listed stocks.

In addition to market-wide NBBO, we look at the Nasdaq exchange in particular and construct its own BBO, again, based on SIP Timestamp ("SIP-Nasdaq-BBO") and Participant Timestamp ("Participant-Nasdaq-BBO") respectively. Recall that both the Nasdaq-SIP and the Nasdaq exchange are located at Carteret, NJ. So for Nasdaq-listed stocks, SIP feed reported from the Nasdaq exchange suffers barely no traveling latency but mainly processing latency at the Nasdaq-SIP. As a result, after the Nasdaq-SIP upgrade, which reduces its processing latency by over 90% to around 20 microseconds, we should see even fewer and shorter dislocations between the two Nasdaq-BBOs in Nasdaq-listed stocks *relative to* NYSE-listed stocks.

To visualize the impact of the Nasdaq-SIP upgrade on NBBO/Nasdaq-BBO dislocations, we compute, for each stock-day combination in our matched sample, two statistics: the number of and the median duration of dislocations for Nasdaq-listed versus NYSE-listed stocks. Figure 3 plots the cross-section average of the two statistics. As expected, after the upgrade, Nasdaq-stocks see a significant reduction in the median duration of both NBBO and Nasdaq-BBO dislocations. For example, the median duration of the Nasdaq-BBO dislocations drops from 200 microseconds to around only 10 microseconds, which implies that for traders who use SIP feeds for their trading in Nasdaq-listed stocks on the Nasdaq exchange, prices they receive become much more up-to-date. As for the number of NBBO and Nasdaq-BBO dislocations, we see a similar pattern: both drop more for Nasdaq-stocks *relative to* NYSE-stocks after the upgrade. In addition, consistent with our expectation, the drop in Nasdaq-BBO dislocations is more pronounced. Quantitatively, the number of NBBO (Nasdaq-BBO) dislocations for Nasdaq-stocks drops by more than 1013 (1354)

³⁶As discussed in detail on page 14, Participant Time in the SIP feeds represents the time when a quote update or trade is registered at the exchange's matching engine and thus suffers no processing latency at and traveling latency from/to the SIP.

³⁷It should be noted that the dislocations we identify here do not necessarily represent true ones experienced by a SIP subscriber compared with a direct feed subscriber. First, even for a direct feed subscriber, she will experience traveling latency and processing latency, although they can be smaller due to better network (e.g, microwave) and better hardware. Second, the actual latency a SIP subscriber experiences depends on its own location, which might be far away from SIPs, making SIP prices she receives even more stale.

than NYSE-stocks, or roughly 18.72% (28.43%) relative to the mean level of Nasdaq-stocks before the upgrade.

Summing up, the foregoing results of the NBBO and Nasdaq-BBO dislocations show that after the Nasdaq-SIP upgrade, SIP prices for Nasdaq-listed stocks, especially on the Nasdaq exchange, experience much less latency and become more reliable relative to NYSE-listed stocks. Whether such an improvement in the SIP prices can translate to better market liquidity or have a significant impact on market participants' trading behavior remains an empirical question we tackle below.

4.1.2 Market-wide impact

Now we examine the impact of the Nasdaq-SIP upgrade on market-wide liquidity and trading by estimating the DiD regression specified in Equation 4. As detailed in Section 3.3.1, the upgrade only reduces the SIP latency of Nasdaq-listed stocks, but not NYSE-listed ones. So if there is any effect of the upgrade, we should see our liquidity and trading variables change significantly for Nasdaq-listed stocks *relative to* NYSE-listed ones.

We first report the estimation results for liquidity variables in Table 8a. It shows that the coefficient on the interaction term *After* × *NasdaqStock*, which captures the DiD effect, is statistically significant for relative quote spread (*RQS*), relative effective spread (*RES*) and relative price impact (*RPI*). Specifically, the three spread measures increase by 0.54, 0.15 and 0.13 basis points respectively for Nasdaq-listed stocks *relative to* NYSE-listed stocks. In terms of economic magnitudes, the increases are about 6.32% ($\approx 0.54 \div 8.55$), 4.76% ($\approx 0.15 \div 3.15$) and 5.53% ($\approx 0.13 \div 2.35$) compared to their unconditional means across all sample stocks and thus relatively small. For the rest of our liquidity measures, the results are insignificant: we do not see a significant change in the NBBO depth, market volume or price improvement received by retail investors.

We then turn to the estimation results of trading variables reported in Table 8b. We find that the coefficient on the interaction term *After* × *NasdaqStock* is significantly positive for the share of ISO trades (*ISOShr*), although the economic magnitude of it is rather small and about 4.37% (\approx 1.53 ÷ 35.03) to its unconditional mean. ISO orders are used to bypass the Reg-NMS order protection rule to trade behind-the-top depth on a target exchange by sweeping through top-of-book depth across all other exchanges with potentially better prices. Thus, as argued by Chakravarty, Jain, Upson,

Table 8. DiD regression: market-wide impact of the Nasdaq-SIP upgrade. All Nasdaq stocks. This table shows the estimation results from the DiD regression specified below:

$$metric_{i,t} = \alpha_i + \beta After_t + \gamma After_t \times NasdaqStock_i + \epsilon_{i,t}$$

where $metr_{i,t}$ is the liquidity or trading variable of stock *i* on day *t*. α_i is the stock fixed effects. After_{*i*,t} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. NasdaqStock_{*i*,t} is dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016.

(a) Liquidity metrics. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact, and relative realized spread respectively and are all in basis point. *Depth* is the NBBO depth in thousand dollars. *Vlm* is trading volume in million dollars. *PrcImp* is share-size weighted average price improvements received by the retail trades in cents per one-hundred shares.

	RQS	RES	RPI	RRS	Depth	Vlm	PrcImp
After	0.63***	0.11***	0.14***	-0.03	-8.97**	25.93***	0.59***
	(0.07)	(0.03)	(0.02)	(0.03)	(3.72)	(2.94)	(0.07)
After x NasdaqStock	0.54***	0.15**	0.13***	0.02	-7.51	-4.02	-0.02
	(0.20)	(0.06)	(0.04)	(0.05)	(13.64)	(6.25)	(0.10)
R ² (%)	2.13	0.27	0.65	0.00	0.21	0.95	0.89
Ν	42900	42900	42900	42900	42900	42900	42900
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes

(b) Trading metrics. *ISOShr* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. *#Run/Vlm* is the number of strategic runs per million dollar trading volume.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade	#Run/Vlm
After	-2.16***	0.09	-8.32***	-14.21***	-2.45***
	(0.16)	(0.10)	(0.44)	(0.98)	(0.42)
After x NasdaqStock	1.53***	0.18	3.09***	6.13***	0.08
	(0.24)	(0.14)	(0.53)	(1.08)	(0.61)
R ² (%)	1.80	0.12	16.16	12.88	1.47
Ν	42900	42900	42900	42900	42900
Stock F.E.	Yes	Yes	Yes	Yes	Yes

and Wood (2012), they are mainly used by informed institutional traders to increase execution speed and capture larger depth. So faster SIP feeds with more up-to-date price information might encourage institutional traders who use them as market data input to submit more ISO orders.

As for odd-lot trades, nowadays they are commonly used by high frequency traders (HFTs) or algorithmic traders (ATs) in strategies such as pinging and order splitting. For example, using the Nasdaq HFT dataset, O'Hara, Yao, and Ye (2014) show that odd lots are more likely to occur when trades are initiated by HFTs. So if faster SIP feeds induce more HFT or overall AT activity, odd-lot share might increase after the upgrade. However, we do not see such an effect. One possible explanation is that SIP feeds do not include odd-lot quotes (see Section 2 on differences between SIP feeds and direct feeds). Therefore, it remains hard for non-HFT ATs who only subscribe to SIP feeds to access odd-lot liquidity even if they become faster. As for HFTs, they shouldn't be affected by faster SIP feeds as they use direct feeds from exchanges as their market data input.

Last, the coefficient on the interaction term is significantly positive for our two AT proxies, *Cancel/Trade* and *Order/Trade*, which indicates that, after the upgrade, there is an increase in overall AT activity for Nasdaq-listed stocks *relative to* NYSE-listed stocks. The economic magnitudes are relatively large: the DiD increases in the two AT proxies are 3.09 and 6.13 respectively, about 12.99% (\approx 3.09 ÷ 23.79) and 16.11% (\approx 6.13 ÷ 38.04) relative to their unconditional means. In contrast, the coefficient for our HFT proxy, *#Run/Vlm*, is statistically insignificant, showing that HFT activity does not change after the upgrade.

To understand the result, it is helpful to stress the difference between the two AT proxies and the HFT proxy: while the former by construction measures the overall level of message traffic, the latter is intended to capture specific HFT strategies which require high-frequency quote revisions in order to react to new information and avoid adverse-selection. So if the coefficient is only statistically significant for the two ATs proxies but not for the HFT proxy, it must be the case that other AT activities unrelated to HFT pick up after the upgrade. For example, one source of AT activity can come from execution algorithms (EAs) which are shown to be dominantly on the passive side of a trade. In one recent study, Beason and Wahal (2021) analyze a large dataset of 2.3 million parent orders executed by institutional investors and find that the dominant order type is limit orders (81.5%) As market condition changes, EAs with algorithms targeting at a certain execution rate (e.g., volume-weighted average price, VWAP, or time-weighted average price, TWAP) will frequently re-price their resting limit orders, i.e., adjust their aggressiveness, in order to raise or lower its execution speed.

Another source of quoting activity, albeit malicious, can result from the quote-stuffing strategy used by ATs or HFTs. For example, Ye, Yao, and Gai (2013) document excessive co-movement between the message flow of two stocks belong to the same Nasdaq-SIP dissemination channel, showing evidence of quote stuffing targeted specifically at SIP feeds. In addition, Egginton, Van

Ness, and Van Ness (2016) identify quote stuffing as intense episodic spikes in quoting activity and find it pervasive, affecting over 74.7% of US-listed stocks during the trading days of 2010. As the upgrade increases the capacity of the Nasdaq-SIP, ill-intended traders will have to send more redundant messages (e.g., quote updates at price levels far from the NBBO) in order to implement the quote stuffing strategy, increasing the quote-to-trade ratio. However, we caution that there is no direct evidence proving the existence of such practice.

Taking stock, we find that faster SIP feeds lead to a mild deterioration of market liquidity in terms of higher spreads and larger price impacts, which might result from more trading from informed institutional traders as we see an elevated level of non-HFT AT activity.

4.1.3 Exchange-specific impact

In the previous section, we examine the potential impact of the Nasdaq-SIP upgrade on *market-wide* liquidity and trading. Now we further look at whether the impact, if any, is more pronounced on the *Nasdaq exchange versus other exchanges*. As we illustrate it in Section 3.3.1, both Nasdaq exchanges and the Nasdaq-SIP are located at Carteret, NJ. The geographical proximity of the two means that messages from Nasdaq exchanges suffer the least traveling latency to the Nasdaq-SIP and thus processing latency at the SIP makes up the largest component of total latency. So after the Nasdaq-SIP upgrade significantly reduces its processing latency, total latency of messages from Nasdaq exchanges should decrease the most relatively. Our empirical strategy then is to run the triple-DiD regression specified in Equation 5 with exchange-specific metrics as dependent variables.

As before, we first look at liquidity variables and report their triple-DiD estimation results in Table 9a. We find that the coefficient on the triple interaction term *After* × *NasdaqStock* × *NasdaqVenue*, which captures the triple-DiD effect, is positively significant for *RES*, *RPI* and *Vlm*. In the DiD results above, we document that both *RES* and *RPI* increase for Nasdaq-listed stocks *relative to* NYSE-listed ones after the Nasdaq-SIP upgrade. Here the triple-DiD results further suggest that the effect is larger on the Nasdaq exchange *relative to* other exchanges, supporting our hypothesis that faster SIP feeds lead to an increase in informed AT activities. As detailed above, while there is an overall reduction of the SIP latency for Nasdaq-listed stocks, the relative Table 9. Triple DiD regression: exchange-specific impact of SIP upgrade. All Nasdaq stocks. This table shows results from the triple-difference-in-difference regression

$$\begin{split} metric_{i,t,e} &= \alpha_{i,e} + \beta After_t \\ &+ \gamma_1 After_t \times NasdaqStock_i \\ &+ \gamma_2 After_t \times NasdaqVenue_e \\ &+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e} \end{split}$$

where *metric*_{*i*,*t*,*e*} is the liquidity or trading metric for stock *i*, traded on exchange *e* and on day *t*. $\alpha_{i,e}$ controls for the stock-venue fixed effects. *After*_{*i*,*t*} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. *NasdaqStock*_{*i*} is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *NasdaqVenue*_{*i*} is dummy variable that equals one if the metric is computed based on exchange *e*, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016.

(a) Liquidity metrics. All measures have been defined above in Table 8a . But note that now they are measures specific to an exchange. For example, *Depth* of an exchange will be its time-weighted dollar depth at the NBBO.

	RQS	RES	RPI	RRS	Vlm	Depth
After	3.87***	0.15***	0.21***	-0.06***	1.71***	-1.05**
	(0.54)	(0.02)	(0.03)	(0.02)	(0.20)	(0.44)
After x NasdaqStock	0.55	0.09*	0.12*	-0.03	-0.02	-1.16
	(0.88)	(0.05)	(0.07)	(0.04)	(0.49)	(2.01)
After x NasdaqVenue	-3.35^{***}	-0.03^{***}	0.01	-0.03^{**}	2.73***	1.44***
	(0.54)	(0.01)	(0.02)	(0.01)	(0.22)	(0.27)
After x NasdaqStock x NasdaqVenue	-0.25	0.06***	0.09***	-0.03	1.78**	0.56
	(0.80)	(0.01)	(0.03)	(0.03)	(0.74)	(1.27)
R ² (%)	0.94	1.05	0.59	0.07	1.02	0.08
Ν	171600	171597	171597	171597	171597	171600
Stock-Venue F.E.	Yes	Yes	Yes	Yes	Yes	Yes

(b) Trading metrics. All measures have been defined above in Table 8b. But note that now they are measures specific to an exchange. For example, *Cancel/Trade* of an exchange will be the cancel-to-trade ratio on its own venue.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	-1.95***	0.28**	-7.69***	-10.86***
	(0.22)	(0.12)	(0.54)	(0.75)
After x NasdaqStock	1.18***	0.21	0.54	0.07
	(0.31)	(0.16)	(0.78)	(1.16)
After x NasdaqVenue	-0.39^{***}	0.42***	-1.08^{***}	-1.47^{**}
	(0.12)	(0.07)	(0.37)	(0.58)
After x NasdaqStock x NasdaqVenue	0.57***	-0.35^{***}	4.33***	5.85***
	(0.19)	(0.11)	(0.59)	(0.99)
R ² (%)	1.01	0.29	2.26	0.19
Ν	171597	171597	171595	171595
Stock-Venue F.E.	Yes	Yes	Yes	Yes

magnitude is much more significant on Nasdaq exchanges. Thus SIP prices for Nasdaq-listed stocks traded on Nasdaq exchanges become much more up-to-date after the upgrade, encouraging

the trading of informed ATs relying on SIP feeds.

Then we turn to trading variables whose triple-DiD estimation results are reported in Table 9b. It shows that the coefficient on the triple-interaction term is significantly positive for *ISOShr* while negative for *OddlotShr*. Recall that ISO orders are used by traders to trade behind-the-top depth at a target exchange. So if faster SIP feeds make the Nasdaq exchange more likely to be the target exchange of ISO traders, then the share of ISO trades on the venue might increase. As for odd lots, although faster SIP feeds in general might lead to more AT activities and thus more odd lots. It is not clear whether the odd-lot share will increase or decrease more on a particular exchange. As already mentioned above, SIP feeds do not include odd-lot quotes and thus it remains hard for non-HFT ATs relying on SIP feeds to access odd-lot liquidity on a particular exchange. However, it is worth mentioning that we are cautious about the triple-DiD results for ISO and odd-lot shares. In Section A.1 of the appendix, we shorten the sample period to two months and the coefficients turn statistically insignificant.

Last, the coefficient on the triple interaction term for the two algorithmic trading (AT) proxies, *Cancel/Trade* and *Order/Trade*, is significantly positive and has similar magnitude as in the DiD results above. In contrast, coefficient on the interaction term *After* × *NasdaqStock* is insignificant, suggesting that the increased AT activity is driven by a higher level of quoting activity in Nasdaq-listed stocks in particular on the Nasdaq exchange. Recall that the DiD results show that while our two AT proxies significantly increase after the Nasdaq-SIP upgrade, the HFT proxy does not change significantly. We suggest two possible explanations: the "algorithmic trading" channel states that faster SIP feeds attract more non-HFT AT activities (e.g., EAs), resulting in more message traffic; in contrast, the "quote stuffing" channel states that a larger SIP capacity forces HFTs who implement quote stuffing strategy to send more "garbage" messages and thus increase the overall message traffic. The triple-DiD results seem to favor the first, benign channel. If HFTs' goal is to slow down the Nasdaq-SIP, it shouldn't matter to which exchange they send "garbage" messages. As a matter of fact, to avoid violating exchange messaging policy³⁸, it is a better strategy to spread "garbage" messages across different exchanges. So if we see the two AT proxies increase more on the Nasdaq exchange than others, it is more likely to be caused by more normal AT activities.

³⁸For example, NASDAQ has an Excessive Messaging Policy that discourages excessive order activity away from the NBBO. Specifically, member firms that exceed a "Weighted Order-to-Trade Ratio" of 100:1 pay a fee on the orders that cause the firm to exceed the threshold.

Table 10. DiD regression: the impact of SIP-glitch events on market liquidity. This table show regression results from the following difference-in-difference regression:

$$metric_{i,d,t} = \alpha_{i,d} + \beta After_{i,d,t} + \gamma After_{d,t} \times Treated_{i,d} + \epsilon_{i,d,t}.$$

where $metr_{i,d,t}$ is the liquidity or trading metric of stock *i* on event day *d* during the 30-second time interval *t*. $\alpha_{i,d}$ is the stock-day fixed effect. $Treated_{i,d}$ is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock on January 3, 2013 and NYSE-listed stock on October 30, 2014 and August 12, 2019. $After_{d,t}$ is a dummy variable that equals one after the glitch starts, and equals zero otherwise. Standard errors are clustered at the event-stock level. The sample stocks consist of 1200 treated stocks (i.e., affected by the SIP glitch) matched with the same number of control stocks on price, market capitalization, trading volume and industry. The sample period is between 30 minutes before the start of the SIP glitch and the end of it. *RQS*, *RES* and *RRS* stand for relative quoted spread, effective spread and realized spread respectively and are all in basis point. *Vlm* is dollar volume in thousand dollars. *DepthNBBO* is NBBO depth in thousand dollars. *Depth5Lvl* is cumulative depth across five best price levels of the order book.

	RQS	RES	RRS	Vlm	DepthNBBO	Depth5Lvl
After	0.35*	-0.04	-0.02	2.85	-3.16	-35.01***
	(0.20)	(0.05)	(0.08)	(1.84)	(2.13)	(4.77)
After x Treated	0.59**	0.37***	0.96***	-5.93^{***}	-17.91^{***}	-64.85^{***}
	(0.25)	(0.07)	(0.10)	(2.14)	(6.55)	(23.42)
R ² (%)	0.22	0.14	0.31	0.01	0.09	0.77
Ν	235564	129309	129309	235564	235564	235564
Stock-Event F.E.	Yes	Yes	Yes	Yes	Yes	Yes

In summary, the triple-DiD results lend support to the hypothesis that faster SIP feeds lead to an increase in non-HFT AT activity from informed institutional traders, resulting in higher spreads and price impacts. Higher message traffic is not likely to be the result of malicious HFT strategies such as quote-stuffing.

4.2 SIP glitches

In the previous section, we examine the impact of a speed upgrade to the SIP feeds, now we turn to glitch events of SIPs and look at what happens to the market when SIP feeds are not available.

4.2.1 Pooling all three events

To study the impact of SIP glitches on market liquidity, we first run the DiD regression specified in Equation 6 where we pool all three major SIP glitch events detailed in Table 1. The identification strategy is identical to what we use in the analysis of SIP upgrade event. Recall that during all three SIP glitch events, only one of the two SIPs is affected. Thus stocks whose SIP feeds are (not) affected serve naturally as the treatment (control) group. So if SIP glitches have a negative impact

on market liquidity, we should see our liquidity measures worsen more for treated stocks than control stocks.

Table 10 reports the regression results. It shows that stocks with their SIP feeds corrupted or unavailable due to a glitch see their market liquidity worsen by all common measures. Specifically, coefficient on the interaction dummy *After* × *Treated*, which captures the DiD effect, shows that relative quoted spread and relative effective spread of treatment stocks increase by 0.59 and 0.37 basis points more than control stocks respectively. In terms of economic magnitudes, they correspond to about 4.28% ($\approx 0.59 \div 13.80$) and 11.14% ($\approx 0.37 \div 3.32$) relative to their unconditional means.

In addition, SIP glitches have a much more significant and negative impact on trading volume and order-book depth. Specifically, trading volume of treatment stocks falls by nearly 17.81% (\approx -5.93 ÷ 33.30) more than control stocks. As for the two order-book depth measures, they follow a similar pattern: NBBO depth and cumulative depth across five best price levels drop by about 21.20% (\approx -17.91 ÷ 84.50) and 13.80% (\approx -64.85 ÷ 470.26) more for treatment stocks than control stocks respectively.

Note that coefficient on the time dummy *After* is significantly positive for relative quoted spread and negative for cumulative depth across five best prices as well, indicating that market liquidity of control stocks are affected as well. The result is perhaps not surprising as trading in the two matched sample of stocks are correlated, either due to them being in the same market index or in the same industry. So when treatment stocks become illiquid, the illiquidity can quickly spread to the control stocks. Such illiquidity contagion through informationally correlated assets are well modeled in Cespa and Foucault (2014). A more simple explanation is that some HFTs or ATs might cease their market making or other trading activities for all stocks when faced with a market-wide data anomaly.

4.2.2 Zooming in onto the event on January 3, 2013

Then we zoom in onto the SIP glitch event on January 3, 2013. First, Figure 4 plots several liquidity measures around the two glitches and visually show that stocks in the early channel and late channel are affected significantly. Then we run the DiD regressions specified in Equation 7 to

Figure 4. Liquidity metrics around the Nasdaq-SIP glitch event on January 3, 2013. This figure plots several liquidity metrics for three groups of stocks respectively. The "Early" group includes Nasdaq-listed stocks belonging to the data channels which were first hit by the glitch. The "Late" groups includes Nasdaq-listed stocks belonging to the data channels which were hit by the glitch later. The "Normal" group includes a matched sample of NYSE-listed stocks. *RQS, RES* stand for relative quoted spread and relative effective spread in basis point. *DollarVlm* is dollar trading volume in millions. *DepthNBBO* is dollar depth at NBBO in thousands. The first two vertical lines indicate the start of the glitch at the "Early" channels an "Late" channels respectively. The last vertical line indicate the end of the glitch.



Table 11. Difference-in-difference regression: Nasdaq-SIP glitch on January 3, 2013. This table show regression results from the following difference-in-difference regression:

$metric_{i,t} = \alpha_i + \beta Period_{i,t} + \beta_1 Period_{i,t} \times Early Channel_{i,t} + \beta_2 Period_{i,t} \times Late Channel_{i,t} + \epsilon_{i,t}$

where $metr_{i,t}$ is the liquidity or trading metric of stock *i* in time interval *t*. α_i is the stock fixed effects. $Period1_{i,t}$ is dummy variable that equals one between the start of the data outage at early channels and late channels, that is, between 13:33:11 and 13:36:51, and equals zero otherwise. Standard errors are clustered at the stock level. $EarlyChannel_{i,t}$ is dummy variable that equals one if stock *i* belongs to one of the early channels, and equals zero otherwise. $LateChannel_{i,t}$ is dummy variable that equals one if stock *i* belongs to one of the late channels, and equals zero otherwise. Note that we include stocks from all early, late and normal channels. The sample period is between 30 minutes before the outage started at early channels and the start of outage at late channels. All variables have been defined in Table 10

	RQS	RES	RRS	Vlm	DepthNBBO	Depth5Lvl
Period1	0.04	-0.03	0.06	-1.38**	3.35	-7.89
	(0.15)	(0.05)	(0.08)	(0.67)	(4.05)	(6.14)
Period1 x EarlyChannel	0.19	0.31*	0.96***	-2.31^{**}	-8.27	-47.03^{**}
	(0.35)	(0.17)	(0.19)	(1.12)	(5.35)	(22.66)
Period1 x LateChannel	0.43	0.01	-0.02	0.45	-0.12	4.08
	(0.31)	(0.09)	(0.16)	(1.18)	(5.01)	(8.53)
R ² (%)	0.09	0.08	0.31	0.04	0.05	0.61
Ν	80886	21803	21803	80886	80886	80886
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

formally test the DiD effect. Table 11 reports the estimation results. The DiD regression is run over the first period of the glitch, that is, when stocks of the even-numbered channels started to experience a glitch but not yet for stocks in the odd-numbered channels. The coefficient on the interaction term *Period1* × *EarlyChannel* then captures the DiD effect of the glitch affected stocks

(i.e., Nasdaq-listed stocks in the even-numbered channels) relative to unaffected stocks (i.e., both Nasdaq-listed stocks in the odd-numbered channels and NYSE-listed stocks). The results show that market liquidity of the affected stocks worsens: relative effective spread increasing by 0.31 basis point or about 10.80% (\approx 0.31 ÷ 2.87) relative to its unconditional mean; cumulative depth across five best prices drops by 47.03 thousand dollars or about 6.64% relative to its unconditional mean (\approx -47.03 ÷ 707.78). As in the regression where all three SIP glitch events are pooled, trading volume is significantly affected. It falls by 2.31 thousand dollars or about 27.66% (\approx -2.31 ÷ 8.35) relative to the unconditional mean.

In summary, both the regression results from the regression pooling all three SIP glitch events and that on the Nasdaq-SIP glitch event show that when SIP feeds are corrupted or unavailable, market liquidity significantly worsens, especially in terms of trading volume and order-book depth.

5 Conclusion

In this paper we study the role of the consolidated feeds in the U.S. equities market by examining exogenous events when they become faster due to technology upgrades and when they are corrupted or unavailable due to technical glitches. The unique structure of two consolidated feeds, one for Nasdaq-listed stocks and the other for NYSE-listed stocks, allows us to implement a standard difference-in-difference (DiD) analysis based on a matched sample of Nasdaq-listed stocks and NYSE-listed stocks. The results show that faster consolidated feeds have a mild and adverse effect on overall market liquidity, with higher spreads and larger price impacts. The worsening of market liquidity might result from an elevated level of non-HFT AT activity form informed institutional traders. In addition, we document that when the consolidated feeds become corrupted or unavailable due to technical glitches, market liquidity, especially market volume and order-book depth, worsens significantly. The foregoing findings show that the consolidated feeds matter and remain a crucial component of the current market data infrastructure. As the consolidated feeds are the focus of the ongoing market structure reform agenda, more careful studies are needed to better assess the potential impacts of proposed new rules.

A Robustness checks

A.1 Event window

In our baseline regression, we focus on the Nasdaq-SIP speed upgrade on October 24, 2016 and use a four-month window, two months before the event date and two months after. As a robustness check, here we shorten the window length to two months, one month before the event date and one month after, to mitigate the concern that we might capture other confounding events other than the Nasdaq-SIP speed upgrade.

Table A1 reports the estimation results of the DiD regression specified in Equation 4. While most results stay qualitatively the same as the baseline, we would like to mention two noticeable changes. On liquidity, the coefficient on the interaction term turns insignificant for *RES*, indicating that while the relative quoted spread increases, traders adapt to finding better prices than NBBO. For example, traders can move to dark pools or access on-exchange hidden liquidity. Turning to trading variables, only the coefficient for *OddlotShr* changes and turns significant. However, its economic magnitude of 0.32% is rather small.

Table A2 reports the estimation results of the triple DiD regression specified in Equation 5. On liquidity, coefficient for all variables are qualitatively the same as the baseline results. On trading, we have already noted in the main text that the coefficient on *ISOShr* and *OddlotShr* turn insignificant. Other key results for the two AT proxies stay qualitatively the same.

A.2 Fixed effect

In our baseline specification for the triple-DiD regression (Equation 5), we control for stock-venue fixed effects. As a robustness check, here we use an alternative specification where we control for

Table A1. DiD regression: market-wide impact of the Nasdaq-SIP upgrade. All Nasdaq stocks. Two-month window.This table shows the estimation results from the DiD regression specified below:

$$metric_{i,t} = \alpha_i + \beta After_t + \gamma After_t \times NasdaqStock_i + \epsilon_{i,t}$$

where $metr_{i,t}$ is the liquidity or trading variable of stock *i* on day *t*. α_i is the stock fixed effects. After_{*i*,t} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. NasdaqStock_{*i*,t} is dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from September 26 to November 18, 2016.

(a) Liquidity metrics. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact, and relative realized spread respectively and are all in basis point. *Depth* is the NBBO depth in thousand dollars. *Vlm* is trading volume in million dollars. *PrcImp* is share-size weighted average price improvements received by the retail trades in cents per one-hundred shares.

	RQS	RES	RPI	RRS	Depth	Vlm	PrcImp
After	0.69***	0.18***	0.28***	-0.10***	-6.32***	31.39***	-0.06
	(0.07)	(0.05)	(0.03)	(0.04)	(2.25)	(3.69)	(0.08)
After x NasdaqStock	0.41**	0.11	0.09*	0.02	-1.49	-0.24	0.14
	(0.20)	(0.08)	(0.05)	(0.06)	(7.70)	(7.74)	(0.11)
R ² (%)	2.12	0.32	1.35	0.06	0.11	1.45	0.01
Ν	21888	21888	21888	21888	21888	21888	21888
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes

(b) Trading metrics. *ISOShr* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. *#Run/Vlm* is the number of strategic runs scaled by trading volume.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade	#Run/Vlm
After	-1.69***	-0.56***	-6.37***	-11.54^{***}	-1.72***
	(0.18)	(0.10)	(0.42)	(0.94)	(0.36)
After x NasdaqStock	1.74***	0.32**	2.62***	5.42***	0.04
	(0.25)	(0.14)	(0.50)	(1.04)	(0.55)
R ² (%)	1.07	0.60	10.64	9.45	0.83
Ν	21888	21888	21888	21888	21888
Stock F.E.	Yes	Yes	Yes	Yes	Yes

the stock-date fixed effects:

$$metric_{i,t,e} = \alpha_{i,t} + \beta Nasdaq Venue_{e}$$

 $+ \gamma_1 NasdaqStock_i \times NasdaqVenue_e$ (8)

 $+ \gamma_2 A fter_t \times Nasdaq Venue_e$

+ $\gamma_3 A fter_t \times Nasdaq Stock_i \times Nasdaq Venue + \epsilon_{i,t,e}$

Table A2. Triple DiD regression: exchange-specific impact of SIP upgrade. All Nasdaq stocks. Two-month window.This table shows results from the triple-difference-in-difference regression

$$\begin{split} metric_{i,t,e} &= \alpha_{i,e} + \beta After_t \\ &+ \gamma_1 After_t \times NasdaqStock_i \\ &+ \gamma_2 After_t \times NasdaqVenue_e \\ &+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e} \end{split}$$

where *metric*_{*i*,*t*,*e*} is the liquidity or trading metric for stock *i*, traded on exchange *e* and on day *t*. $\alpha_{i,e}$ controls for the stock-venue fixed effects. *After*_{*i*,*t*} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. *NasdaqStock*_{*i*} is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *NasdaqVenue*_{*i*} is dummy variable that equals one if the metric is computed based on exchange *e*, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from September 26 to November 18, 2016.

(a) Liquidity metrics. All measures have been defined above in Table 8a. But note that now they are measures specific to an exchange. For example, *Depth* of an exchange will be its time-weighted dollar depth at the NBBO.

	RQS	RES	RPI	RRS	Vlm	Depth
After	3.68***	0.21***	0.35***	-0.14^{***}	2.20***	-0.54*
	(0.52)	(0.02)	(0.04)	(0.03)	(0.25)	(0.29)
After x NasdaqStock	-0.20	0.05	0.08	-0.04	0.54	-0.27
	(0.78)	(0.05)	(0.06)	(0.04)	(0.64)	(1.03)
After x NasdaqVenue	-2.88^{***}	-0.02^{***}	0.04**	-0.05^{***}	2.96***	0.71**
	(0.46)	(0.01)	(0.02)	(0.02)	(0.28)	(0.29)
After x NasdaqStock x NasdaqVenue	0.32	0.06***	0.09***	-0.03	2.76***	1.85
	(0.67)	(0.02)	(0.03)	(0.03)	(0.97)	(1.22)
R ² (%)	0.69	1.74	1.21	0.28	1.62	0.04
Ν	87552	87552	87552	87552	87552	87552
Stock-Venue F.E.	Yes	Yes	Yes	Yes	Yes	Yes

(b) Trading metrics. All measures have been defined above in Table 8b. But note that now they are measures specific to an exchange. For example, *Cancel/Trade* of an exchange will be the cancel-to-trade ratio on its own venue.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	-1.97***	-0.52***	-6.74^{***}	-9.95***
	(0.22)	(0.12)	(0.63)	(0.92)
After x NasdaqStock	1.92***	0.33**	0.98	0.43
	(0.31)	(0.16)	(0.96)	(1.81)
After x NasdaqVenue	-0.45^{***}	-0.10	-0.68	-1.03
	(0.13)	(0.08)	(0.47)	(0.71)
After x NasdaqStock x NasdaqVenue	-0.01	-0.09	3.43***	5.41***
	(0.20)	(0.11)	(0.80)	(1.65)
R ² (%)	1.03	0.27	1.58	0.09
Ν	87552	87552	87551	87551
Stock-Venue F.E.	Yes	Yes	Yes	Yes

where *metric*_{*i*,*t*,*e*} is the liquidity or trading metric for stock *i*, traded on exchange *e* and on day *t*. $\alpha_{i,t}$ controls for the stock-day fixed effects. *After*_{*i*,*t*} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. *NasdaqStock*_i is a dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *NasdaqVenue*_i is dummy variable that equals one if the metric is computed based on exchange e, and equals zero otherwise.

As argued in Brogaard and Brugler (2021), the two types of fixed effects, stock-venue used in the baseline regression and stock-day used here, control for different variations. By including stock-venue fixed effects, we allow for separate intercepts for each stock-venue pair (e.g., Apple traded on Nasdaq), which is important as trading of a particular stock varies significantly across venues, especially between its listing exchange and other non-listing exchanges. For example, trading volume of a stock is on average larger on its listing exchange (e.g., Apple is more traded on Nasdaq versus NYSE). Besides, NYSE-listed stocks are assigned to what's called DMMs (designated market makers), who have mild obligations for maintaining liquidity. In contrast, by including stock-day fixed effects, we control for unobserved effects at the stock-day level which affect trading of the stock across all venues. However, we have to impose an implicit assumption that differences within a stock across venues is constant across stocks (e.g., the effect is the same for Apple traded on Nasdaq vs. NYSE as for Nvidia traded on Nasdaq vs. NYSE).

Table A3 reports the estimation results based on the alternative specification with stock-day fixed effects as in Equation 8. It shows that coefficients on the triple interaction term are very similar to those based on the baseline specification.

Table A3. Triple-DiD regression: exchange-specific impact of SIP upgrade. Stock-day fixed effects. This table shows results from the triple-DiD regression

$$\begin{split} metric_{i,t,e} &= \alpha_{i,t} + \beta NasdaqVenue_e \\ &+ \gamma_1 NasdaqStock_i \times NasdaqVenue_e \\ &+ \gamma_2 After_t \times NasdaqVenue_e \\ &+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e} \end{split}$$

where *metric*_{*i*,*t*,*e*} is the liquidity or trading metric for stock *i*, traded on exchange *e* and on day *t*. $\alpha_{i,t}$ controls for the stock-day fixed effects. *After*_{*i*,*t*} is dummy variable that equals one after October 24, 2016, and equals zero otherwise. *NasdaqStock*_{*i*} is dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *NasdaqVenue*_{*i*} is dummy variable that equals one if stock *i* is a Nasdaq-listed stock, and equals zero otherwise. *Standard errors are clustered at the stock level*. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016.

(a) Liquidity metrics. All measures have been defined above in Table 8a. But note that now they are measures specific to an exchange. For example, *Depth* of an exchange will be its time-weighted dollar depth at the NBBO.

	RQS	RES	RPI	RRS	Vlm	Depth
NasdaqVenue	-12.96***	0.00	-0.07***	0.06***	5.27***	10.65***
	(1.04)	(0.01)	(0.02)	(0.01)	(0.34)	(0.60)
NasdaqStock x NasdaqVenue	-2.46	0.11***	0.60***	-0.49^{***}	9.48***	16.39***
	(1.66)	(0.01)	(0.04)	(0.03)	(1.50)	(2.31)
After x NasdaqVenue	-3.35***	-0.03^{***}	0.01	-0.03^{**}	2.73***	1.44***
	(0.54)	(0.01)	(0.02)	(0.01)	(0.22)	(0.27)
After x NasdaqStock x NasdaqVenue	-0.25	0.06***	0.09***	-0.03	1.78**	0.56
	(0.80)	(0.01)	(0.03)	(0.03)	(0.74)	(1.27)
R ² (%)	7.06	0.49	1.93	1.38	18.60	11.42
Ν	171600	171597	171597	171597	171597	171600
Stock-Date F.E.	Yes	Yes	Yes	Yes	Yes	Yes

(b) **Trading metrics.** All measures have been defined above in Table 8b. But note that now they are measures specific to an exchange. For example, *Cancel/Trade* of an exchange will be the cancel-to-trade ratio on its own venue.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
NasdaqVenue	1.01***	2.81***	0.26	3.11***
	(0.11)	(0.13)	(0.52)	(0.78)
NasdaqStock x NasdaqVenue	-2.48^{***}	0.46**	-11.77^{***}	-17.67^{***}
	(0.22)	(0.22)	(1.09)	(2.87)
After x NasdaqVenue	-0.39***	0.42***	-1.09^{***}	-1.48^{**}
	(0.12)	(0.07)	(0.37)	(0.58)
After x NasdaqStock x NasdaqVenue	0.57***	-0.35^{***}	4.36***	5.96***
	(0.19)	(0.11)	(0.61)	(1.06)
R ² (%)	1.01	17.28	1.45	0.13
Ν	171597	171597	171595	171595
Stock-Date F.E.	Yes	Yes	Yes	Yes

B Visual evidence for parallel trend assumptions

Figure A1 plots the market-wide liquidity and trading metrics used in the DiD regression and presents visual evidence for the parallel trend assumptions. Figure A2 plots the exchange-specific liquidity and trading metrics used in the triple DiD regression and presents visual evidence for the parallel trend assumptions.

C SIP glitch events

To illustrate the three SIP technical glitch events, we plot, for each event, the number of trades and quote updates from direct feeds and consolidated feeds by exchange. Specifically, Figure A3 plots the Nasdaq-SIP glitch event on January 3, 2013; Figure A4 plots the same metrics for the NYSE-SIP glitch event on October 30, 2014; Figure A5 plots the same metrics for the NYSE-SIP glitch event on August 22, 2019.

Figure A1. Daily liquidity and trading metrics for NYSE-listed stocks versus Nasdaq-listed stocks around the Nasdaq-SIP speed upgrade on October 24, 2016. This figure plots the daily time-series of several liquidity and trading metrics for Nasdaq-listed stocks and a matched sample of NYSE-listed stocks. The vertical line represents the speed upgrade to the Nasdaq-SIP on October, 24, 2016. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact and relative realized spread respectively and are all in basis point. *Depth* is NBBO depth in thousand dollars. *PrcImp* is the share-size weighted average price improvements received by the retail trades and in cents per hundred shares. *Vlm* is trading volume in million dollars. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *ISOShare* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. *#Run/Vlm* is the number of strategic runs per million dollar trading volume. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume and industry. The sample period is from August 29 to December 16, 2016.



Figure A2. Daily liquidity and trading metrics on the Nasdaq exchange versus NYSE Arca and Bats around the Nasdaq-SIP speed upgrade on October 24, 2016. This figure plots the daily time-series of several liquidity and trading metrics on the Nasdaq exchange and NYSE Arca and Bats. The vertical line represents the speed upgrade to the Nasdaq-SIP on October, 24, 2016. The sample period is from September 26 to December 1, 2016. *RQS, RES* stand for relative quoted spread and relative effective spread in basis point.*Depth* is dollar depth at NBBO in thousands. *Vlm* is dollar trading volume in millions. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. *PrcImpShr* is the trade volume that receives price improvement as a fraction of total dark trading volume. *ISOShare* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume.



Figure A3. Nasdaq-SIP glitch event on January 3, 2013. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidate feeds around the Nasdaq-SIP glitch event on January 3, 2013.

(a) Number of quotes.



(b) Number of trades.





Figure A4. Nasdaq-SIP glitch event on October 30, 2014. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidate feeds around the Nasdaq-SIP glitch event on October 30, 2014.



(a) Number of quotes.

(b) Number of trades.





Figure A5. Nasdaq-SIP glitch event on August 12, 2019. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidate feeds around the Nasdaq-SIP glitch event on August 12, 2019.



(a) Number of quotes.

(b) Number of trades.



References

- Aldrich, Eric M., Joseph Grundfest, and Gregory Laughlin (2017). "The Flash Crash: A New Deconstruction".
- Bartlett, Robert P. and Justin McCrary (2019). "How rigged are stock markets? Evidence from microsecond timestamps". In: *Journal of Financial Markets* 45, pp. 37–60.
- Battalio, Robert H., Shane A. Corwin, and Robert Jennings (2016). "Unrecognized odd lot liquidity supply". In: *Journal of Trading*, pp. 35–41.
- Beason, Tyler and Sunil Wahal (2021). "The Anatomy of Trading Algorithms". In: SSRN Electronic Journal.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas (2015). "Equilibrium fast trading". In: *Journal of Financial Economics* 116.2, pp. 292–313.
- Boehmer, Ekkehart, Charles M. Jones, Xiaoyan Zhang, and Xinran Zhang (2021). "Tracking Retail Investor Activity". In: *The Journal of Finance*, jofi.13033.
- Brogaard, Jonathan and James Brugler (2021). "Competition and Exchange Data Fees".
- Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan (2015). "Trading Fast and Slow: Colocation and Liquidity". In: *Review of Financial Studies* 28.12, pp. 3407–3443.
- Brogaard, Jonathan, Matthew C. Ringgenberg, and Dominik Rösch (2020). "Does Floor Trading Matter?"
- Budish, Eric B., Peter Cramton, and John J. Shim (2015). "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response". In: *The Quarterly Journal of Economics* 130.4, pp. 1547–1621.
- Cespa, Giovanni and Thierry Foucault (2014). "Illiquidity contagion and liquidity crashes". In: *Review of Financial Studies* 27.6, pp. 1615–1660.
- CFTC and SEC (2010). Findings Regarding the Market Events of May 6, 2010. Tech. rep.
- Chakravarty, Sugato, Pankaj K. Jain, James Upson, and Robert Wood (2012). "Clean sweep: Informed trading through intermarket sweep orders". In: *Journal of Financial and Quantitative Analysis* 47.2, pp. 415–435.
- Ding, Shengwei, John Hanna, and Terrence J. Hendershott (2014). "How slow is the NBBO? A comparison with direct exchange feeds". In: *Financial Review* 49.2, pp. 313–332.

Egginton, Jared F., Bonnie F. Van Ness, and Robert A. Van Ness (2016). "Quote Stuffing". In: *Financial Management* 45.3, pp. 583–608.

Ernst, Thomas, Jonathan Sokobin, and Chester Spatt (2021). "The Value of Off-Exchange Data".

- Foucault, Thierry, Johan Hombert, and Ioanid Roşu (2016). "News Trading and Speed". In: *Journal of Finance* 71.1, pp. 335–382.
- Hagströmer, Björn and Lars Nordén (2013). "The diversity of high-frequency traders". In: *Journal* of *Financial Markets* 16.4, pp. 741–770.
- Hasbrouck, Joel (2019). "Price Discovery in High Resolution". In: *Journal of Financial Econometrics*, pp. 1–36.
- Hasbrouck, Joel and Gideon Saar (2013). "Low-latency trading". In: *Journal of Financial Markets* 16.4, pp. 646–679.
- Hendershott, Terrence J., Marc Rysman, and Rainer Schwabe (2020). "Stock Exchanges as Platforms for Data and Trading".
- Hoffmann, Peter (2014). "A dynamic limit order market with fast and slow traders". In: *Journal of Financial Economics* 113.1, pp. 156–169.
- IEX (2019). The Cost Of Exchange Services. Tech. rep. January.
- Jovanovic, Boyan and Albert J. Menkveld (2016). "Middlemen in Limit-Order Markets".
- Lee, Charles M. C. and Mark J. Ready (1991). "Inferring Trade Direction from Intraday Data". In: *The Journal of Finance* 46.2, pp. 733–746.
- Menkveld, Albert J. (2016). "The Economics of High-Frequency Trading: Taking Stock". In: *Annual Review of Financial Economics* 8.1, pp. 1–24.
- (2018). "High-frequency trading as viewed through an electron microscope". In: *Financial Analysts Journal* 74.2, pp. 24–31.
- Menkveld, Albert J. and Marius Zoican (2016). "Need for Speed ? Exchange Latency and Liquidity". In: *Review of Financial Studies* 30.4, pp. 1188–1228.
- O'Hara, Maureen, Chen Yao, and Mao Ye (2014). "What's not there: Odd lots and market data". In: *Journal of Finance* 69.5, pp. 2199–2236.
- Shkilko, Andriy and Konstantin Sokolov (2020). "Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity, and Trading Costs". In: *Journal of Finance* 75.6, pp. 2899–2927.

- Yao, Chen and Mao Ye (2018). "Why trading speed matters: A tale of queue rationing under price controls". In: *Review of Financial Studies* 31.6, pp. 2157–2183.
- Ye, Mao (2012). "Price Manipulation, Price Discovery and Transaction Costs in the Crossing Network".
- Ye, Mao, Chen Yao, and Jiading Gai (2013). "The Externality of High Frequency Trading".